EXAMTOPICS

-Expert Verified, Online, **Free**.



CERTIFICATION TEST

-CertificationTest.net - Cheap & Quality Resources With Best Support

Question #1 *Topic 1*

Snowflake provides a mechanism for its customers to override its natural clustering algorithms. This method is:

A. Micro-partitions

B. Clustering keys

C. Key partitions

D. Clustered partitions

**Suggested Answer:** *B*

*Community vote distribution*

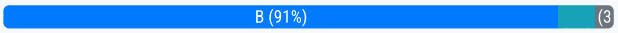


**aikazch** Most Recent 4 months, 3 weeks ago

B Clustering keys

upvoted 1 times

**Nari123456** 10 months, 3 weeks ago

B Clustering keys

upvoted 1 times

**nitesh1206** 1 year, 1 month ago

create a Cluster key

upvoted 1 times

**baskis** 1 year, 3 months ago

Create customer cluster keys

upvoted 1 times

**souhaaaaaaa** 1 year, 6 months ago

.Clustering keys correct

upvoted 1 times

**zonde123** 1 year, 6 months ago

Clustering keys correct

upvoted 1 times

**DanRaja** 1 year, 6 months ago

Snowflake automatically organizes data into micro-partitions, which are contiguous units of storage optimized for query performance. However, users can override Snowflake's natural clustering algorithms by defining clustering keys. Clustering keys allow users to specify one or more columns to reorganize data within micro-partitions, optimizing query performance and enabling efficient pruning of irrelevant micro-partitions during queries

upvoted 4 times

**Yayi123** 1 year, 7 months ago

Clustering keys

upvoted 1 times

**oragenashok** 1 year, 7 months ago

Snowflake provided micro partitions to store and retrieve data faster than native algorithms that other database used.

upvoted 1 times

**Atuu3011** 1 year, 8 months ago

Clustering keys are used to override the default cluster keys.

upvoted 2 times

**RatishKumar** 1 year, 9 months ago

Option B - Clustering Key is the correct answer.

upvoted 1 times

**tharangaweer** 1 year, 9 months ago

Clustering key

upvoted 1 times

**NeelimaKambam** 1 year, 11 months ago

clustering keys

upvoted 1 times

**Djama** 1 year, 11 months ago

clustering keys

upvoted 1 times

**Snowuser2134** 2 years, 2 months ago

correct answer B

upvoted 1 times

**sayan_data_god** 2 years, 3 months ago

clustering keys

upvoted 1 times

**nani_1998** 2 years, 7 months ago

**Selected Answer: B**

Clustering Key

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B Clustering Keys

upvoted 1 times

Question #2 *Topic 1*

Which of the following are valid Snowflake Virtual Warehouse Scaling Policies? (Choose two.)

A. Custom

B. Economy

C. Optimized

D. Standard

**Suggested Answer:** *B, D*

*Community vote distribution*

BD (100%)

**bpxgeek** Most Recent 10 months ago

A &D is the correct answer

upvoted 1 times

**nitesh1206** 1 year, 1 month ago

Standard and economy

upvoted 1 times

**DanRaja** 1 year, 6 months ago

Economy : Prioritizes cost efficiency by conserving credits. A new cluster is provisioned only if the existing clusters are expected to remain busy for at least 6 minutes. Clusters are terminated after 5–6 minutes of inactivity checks. Standard: Minimizes query queuing to prioritize performance. New clusters are added immediately when queries are queued or existing clusters are overloaded. Clusters are terminated after 2–3 minutes of inactivity checks.

upvoted 2 times

**oragenashok** 1 year, 7 months ago

Standard and Economy are the two scaling capacity that snowflake provides.

upvoted 2 times

**Atuu3011** 1 year, 8 months ago

Economy and standard are two valid scaling WH

upvoted 2 times

**Faissal** 1 year, 9 months ago

Economy and Standard

upvoted 1 times

**tharangaweer** 1 year, 9 months ago

Economy and Standard are the sf valid scaling

upvoted 1 times

**m_sflk** 2 years ago

B and D

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B and D

upvoted 1 times

**nani_1998** 2 years, 7 months ago

Economy & Standard

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

B & D

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B-Economy and D-Standard

upvoted 1 times

**Feb01** 2 years, 9 months ago

B-Economy & D-Standard

upvoted 1 times

**KHarish** 2 years, 11 months ago

B and D

upvoted 1 times

**Nagaraju_S** 2 years, 12 months ago

option b and d

upvoted 1 times

**Manish5031994** 3 years, 1 month ago

End of discussion answer Economy and Standard.

upvoted 1 times

**MikeTQ** 3 years ago

Economy and standard

upvoted 1 times

**galicneverdie** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/warehouses-multicluster#setting-the-scaling-policy-for-a-multi-cluster-warehouse

upvoted 2 times

**Shanvi_27** 3 years, 1 month ago

the best answers are : B and D

upvoted 1 times

Question #3 *Topic 1*

True or False: A single database can exist in more than one Snowflake account.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Sundarr** Highly Voted 4 years, 9 months ago

A database is a logical grouping of schemas. Each database belongs to a single Snowflake account. A schema is a logical grouping of database objects (tables, views, etc.). Each schema belongs to a single database.

upvoted 14 times

**Rad1941** Highly Voted 2 years, 1 month ago

Joe Kokomo has a Snowflake account. Joe has a car dealership in New York. Joe's database is named JOE_DB. pronounded JODIE-B. Joe Piscapo has a Snowflake account. Joe has a car dealership in Nebraska. Joes' database is named JOE_DB. pronounced JODIE-B. These databases contain the exact same data. 3000 rows of data purchased from Car Dealers Consortium in Denver CO. Recap: Same name. Same physical properties. Same content. Different accounts. The question is 'can the same database exist in more than1 Snowflake account'. 100% of you say no.

upvoted 5 times

**nitesh1206** Most Recent 1 year, 1 month ago

false, a database can be replicated but it won't be same

upvoted 1 times

**84d4057** 1 year, 5 months ago

it is false

upvoted 1 times

**dbairfan** 1 year, 6 months ago

Database and account have one to one mapping.

upvoted 2 times

**Yayi123** 1 year, 7 months ago

No a single db is only one SF account as well

upvoted 1 times

**oragenashok** 1 year, 7 months ago

B . False

upvoted 1 times

**tharangaweer** 1 year, 9 months ago

B.False

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

b False

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

b.False

upvoted 1 times

**_yyukta** 2 years, 8 months ago

b.False

upvoted 1 times

**Feb01** 2 years, 9 months ago

B-False

upvoted 1 times

**KHarish** 2 years, 11 months ago

Account can have multiple Databases. But Database can be associated to that single snowflake account. B-False

upvoted 2 times

**Shanvi_27** 3 years, 1 month ago

the best answer is False

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Ravi9292** 3 years, 4 months ago

False,one account can have many databases .But one database should belong to one account only.

upvoted 2 times

**instride** 3 years, 5 months ago

False; Each DB belongs only to a single snowflake account

upvoted 2 times

Question #4 *Topic 1*

Which of the following roles is recommended to be used to create and manage users and roles?

A. SYSADMIN

B. SECURITYADMIN

C. PUBLIC

D. ACCOUNTADMIN

**Suggested Answer:** *B*

*Community vote distribution*

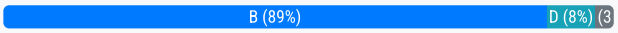


**blazingliverbird** Highly Voted 4 years, 9 months ago

I think the answer to this should be USERADMIN which is not in the options. So the next best answer will be SECURITYADMINSource: https://docs.snowflake.com/en/user-guide/security-access-control-considerations The user administrator (USERADMIN) role includes the privileges to create and manage users and roles (assuming ownership of those roles or users has not been transferred to another role). The security administrator (SECURITYADMIN system-defined) role includes the global MANAGE GRANTS privilege to grant or revoke privileges on objects in the account. >>>The USERADMIN role is a child of this role in the default access control hierarchy.

upvoted 24 times

**blazingliverbird** Highly Voted 4 years, 9 months ago

I think the answer to this should be USERADMIN which is not in the options. So the next best answer will be SECURITYADMIN

upvoted 23 times

**BigDataBB** Most Recent 4 months, 1 week ago

Source: https://docs.snowflake.com/en/user-guide/security-access-control-considerations The user administrator (USERADMIN) role includes the privileges to create and manage users and roles (assuming ownership of those roles or users has not been transferred to another role). The security administrator (SECURITYADMIN system-defined) role includes the global MANAGE GRANTS privilege to grant or revoke privileges on objects in the account. >>>The USERADMIN role is a child of this role in the default access control hierarchy.

upvoted 1 times

**c0d3g** 3 years, 8 months ago

Here the answer should be SECURITYADMIN as it is parent of USERADMIN (which is the correct answer) More details: https://docs.snowflake.com/en/user-guide/security-access-control-considerations Attention By default, when your account is provisioned, the first user is assigned the ACCOUNTADMIN role. This user should then create one or more additional users who are assigned the USERADMIN role. All remaining users should be created by the user(s) with the USERADMIN role or another role that is granted the global CREATE USER privilege.

upvoted 10 times

**nitesh1206** 1 year, 1 month ago

Security Admin role

upvoted 1 times

**houssam_ben** 1 year, 4 months ago

the answear is D because the ACCOUNTADMIN in the snowflake documentation its "Role that encapsulates the SYSADMIN and SECURITYADMIN system-defined roles" , also in documentation USERADMIN is "Role that is dedicated to user and role management only."

upvoted 1 times

**tardigradum** 1 year, 5 months ago

The question is flawed. The correct role for creating and managing users and roles in Snowflake is USERADMIN, which is not listed among the answer choices. While SECURITYADMIN can assign roles and manage grants, --->it cannot create users<--- (guys, be careful with this!!!!!). The only listed role that technically can create users is ACCOUNTADMIN, but using it for routine user management is not a best practice—it's overly privileged and should be reserved for account-level administration only. Conclusion: If USERADMIN is not an option, the only technically correct choice is D. ACCOUNTADMIN, even though the question is poorly constructed and encourages bad practice.

upvoted 1 times

**84d4057** 1 year, 5 months ago

answar B

upvoted 1 times

**dbairfan** 1 year, 6 months ago

Securityadmin is recommended for user management.

upvoted 1 times

**Yayi123** 1 year, 7 months ago

SECURITYADMIN

upvoted 1 times

**PrasantSadangi** 1 year, 7 months ago

Security Admin

upvoted 1 times

**tharangaweer** 1 year, 9 months ago

ORGADMIN manages operations at the organizational level ACCOUNTADMIN performs administrative functions SECURITYADMIN creates and manages users and roles, and manages grants USERADMIN creates users and roles SYSADMIN creates and manages databases, schemas, and warehouses PUBLIC is available to all users, but (by default) cannot do anything but log in

upvoted 1 times

**hotflake** 2 years, 3 months ago

TYPES OF ROLES: SYSTEM ROLES: ● ORGADMIN manages operations at the organizational level ● ACCOUNTADMIN performs administrative functions ● SECURITYADMIN creates and manages users and roles, and manages grants ● USERADMIN creates users and roles ● SYSADMIN creates and manages databases, schemas, and warehouses ● PUBLIC is available to all users, but (by default) cannot do anything but log in CUSTOM ROLES: ● Created by the SECURITYADMIN or USERADMIN based on a company's desired hierarchy ● Custom roles generally "roll up" to the SYSADMIN role

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

B.Securityadmin

upvoted 1 times

**pranalig** 2 years, 4 months ago

correct answer:B

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

B. SECURITYADMIN The SECURITYADMIN role is specifically designed to handle security-related tasks, including the creation and management of users and roles. This role has the necessary privileges to manage all aspects of user and role administration, which is essential for maintaining a secure and well-managed Snowflake environment.

upvoted 1 times

**JasMozai** 2 years, 6 months ago

Security admin: creates and manges users and roles, and manages grants

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Answer B

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. SECURITYADMIN

upvoted 1 times

Question #5 *Topic 1*

True or False: Bulk unloading of data from Snowflake supports the use of a SELECT statement.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**nitesh1206** 1 year, 1 month ago

True, Bulk loading support select query

upvoted 3 times

**SahuVijay** 4 months ago

**Selected Answer: A, B, C**

Named stage => @namedstage Table stage => @%tablestage User stage => @~userstage

upvoted 3 times

**dbairfan** 1 year, 6 months ago

True Bulk unload supports select.

upvoted 3 times

**dbairfan** 4 months ago

**Selected Answer: A, B, C**

Named Stage, User Stage, Table Stage

upvoted 2 times

**PrasantSadangi** 1 year, 7 months ago

A - True

upvoted 1 times

**Yayi123** 4 months ago

**Selected Answer: A, B, C**

The stages r 3 types User stage Table stage and named stage..

upvoted 2 times

**tharangaweer** 1 year, 9 months ago

A.true

upvoted 1 times

**4762448** Most Recent 4 months ago

**Selected Answer: D**

Schema Stage is not a Snowflake's internal stage.

upvoted 2 times

**SahuVijay** 1 year, 9 months ago

Unloading allows to store data in multiple files. It also allows to join which is not allowed in loading.

upvoted 1 times

**nitesh1206** 4 months ago

**Selected Answer: A, B, C**

Schema stage doesn't exist

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A.True

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. true

upvoted 1 times

**KHarish** 2 years, 11 months ago

True A

upvoted 1 times

**Shanvi_27** 3 years, 1 month ago

the best answer is : A Snowflake supports specifying a SELECT statement instead of a table in the COPY INTO command. The results of the query are written to one or more files as specified in the command and the file(s) are stored in the specified location (internal or external).

upvoted 1 times

**galicneverdie** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/data-unload-overview#bulk-unloading-using-queries

upvoted 2 times

**asorock** 3 years, 3 months ago

Bullk unload supports "SELECT" statements and the answer is True. See query below COPY INTO @my_stage/my_unloaded_data.csv FROM ( SELECT col1, col2, col3 FROM my_table WHERE condition ) FILE_FORMAT = (FORMAT_NAME = 'CSV' COMPRESSION = 'AUTO');

upvoted 2 times

**VS3497** 3 years, 4 months ago

False Unload into file_name from my_table File_format = CSV

upvoted 1 times

**shubtred** 3 years, 4 months ago

B UNLOAD INTO @stage_name/files FROM my_table FILE_FORMAT = (TYPE = 'CSV');

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 2 times

**Fabmel1** 3 years, 6 months ago

It's a little bit confuse due to BULK keyword, but the correct answer is A.

upvoted 1 times

**cxna** 3 years, 6 months ago

A is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

A is correct

upvoted 1 times

Question #6 *Topic 1*

Select the different types of Internal Stages: (Choose three.)

A. Named Stage

B. User Stage

C. Table Stage

D. Schema Stage

**Suggested Answer:** *A, B, C*

*Community vote distribution*

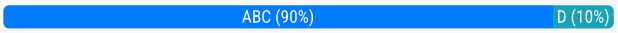


**DataSam24** Most Recent 2 months, 4 weeks ago

User and Table stages are allocated by SNOWFLAKE for each user and table automatically. A named stage is created by a user using command CREATE STAGE mistake

upvoted 1 times

**nitesh1206** 1 year, 1 month ago

Schema stage doesn't exist

upvoted 1 times

**korean1414** 1 year, 3 months ago

Name Stage, User Stage, Table Stage

upvoted 1 times

**dbairfan** 1 year, 6 months ago

Named Stage, User Stage, Table Stage

upvoted 2 times

**Aman7288** 1 year, 7 months ago

Named, User and Table Stage A,B,C

upvoted 1 times

**Yayi123** 1 year, 7 months ago

The stages r 3 types User stage Table stage and named stage..

upvoted 2 times

**4762448** 1 year, 7 months ago

Schema Stage is not a Snowflake's internal stage.

upvoted 2 times

**SahuVijay** 1 year, 9 months ago

Named stage => @namedstage Table stage => @%tablestage User stage => @~userstage

upvoted 3 times

**Mallikharjuna452** 2 years, 3 months ago

A,B,C is correct answer

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

User Table named

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. Named stage B.user Stage C.Table stage

upvoted 1 times

**KHarish** 2 years, 11 months ago

Internal: ABC( Named, Table and User stages)

upvoted 1 times

**Shanvi_27** 3 years, 1 month ago

Snowflake Supports three types of stages. User Stage Table Stage Internal Named Stage

upvoted 1 times

**galicneverdie** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/data-load-local-file-system-create-stage#types-of-internal-stages

upvoted 1 times

**VS3497** 3 years, 4 months ago

User Table named

upvoted 1 times

**HitenPat** 3 years, 4 months ago

ABC is the correct answer

upvoted 1 times

**cxna** 3 years, 6 months ago

**Selected Answer: ABC**

ABC is the correct answer

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

ABC is the correct answer

upvoted 1 times

Question #7 *Topic 1*

True or False: A customer using SnowSQL / native connectors will be unable to also use the Snowflake Web Interface (UI) unless access to the UI is explicitly granted by support.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**nitesh1206** Most Recent 1 year, 1 month ago

B option - False

upvoted 1 times

**Aman7288** 1 year, 7 months ago

B. False

upvoted 1 times

**Yayi123** 1 year, 7 months ago

B :false

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B.False

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

B . False

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B.False

upvoted 1 times

**Shanvi_27** 3 years, 1 month ago

best answer is : False

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**cxna** 3 years, 6 months ago

B is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

B is correct

upvoted 1 times

**Percy2112** 4 years, 2 months ago

No need for extra privileges.

upvoted 2 times

**rox2k** 4 years, 5 months ago

B is correct

upvoted 2 times

**Sundarr** 4 years, 9 months ago

WebUI and SnowSQL are available for all users. Answer is FALSE

upvoted 4 times

**Julien25** 4 years, 9 months ago

WebUI available for all user

upvoted 1 times

Question #8 *Topic 1*

Account-level storage usage can be monitored via:

A. The Snowflake Web Interface (UI) in the Databases section

B. The Snowflake Web Interface (UI) in the Account -> Billing & Usage section

C. The Information Schema -> ACCOUNT_USAGE_HISTORY View

D. The Account Usage Schema -> ACCOUNT_USAGE_METRICS View

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**bharathmv** Highly Voted 4 years, 2 months ago

B is correct. for D it should be TABLE_STORAGE_METRICS not ACCOUNT_USAGE_METRICS. there is no view with account_usage_metrics name. if D has said view name is TABLE_STORAGE_METRICS then its correct

upvoted 10 times

**zonde123** Most Recent 1 year, 6 months ago

B is correct answer

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B. The Snowflake Web Interface (UI) in the Account -> Billing & Usage section

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

B is correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B.The Snowflake Web Interface (UI) in the Account -> Billing & Usage section

upvoted 1 times

**Jules27** 2 years, 10 months ago

I think A is the most correct. https://docs.snowflake.com/en/user-guide/cost-exploring-data-storage B. There is no path: Account -> Billing & Usage. The correct path is Admin -> Cost Management. C. The Information Schema will not contain account usage data. D. There is no ACCOUNT_USAGE_METRICS View.

upvoted 2 times

**KHarish** 2 years, 11 months ago

B is correct

upvoted 1 times

**Shanvi_27** 3 years, 1 month ago

best answer is : The Snowflake Web Interface (UI) in the Account -> Billing & Usage section

upvoted 1 times

**Marge23** 3 years, 1 month ago

Consider A, in Database section with 3 locations. Total Storage cost is the sum of costs associated with: 1) Staged file storage 2) Database table storage 3) Fail-safe and Time Travel storage To use Snowsight to explore storage costs: 1. Navigate to Admin » Usage. 2. Select a warehouse to use to view the usage data. Snowflake recommends using an XS warehouse for this purpose. 3. Select Storage from the Usage Type drop-down. https://docs.snowflake.com/en/user-guide/cost-exploring-data-storage

upvoted 1 times

**Marge23** 3 years, 1 month ago

Consider A, in Database section with 3 locations. Total Storage cost is the sum of costs associated with: 1) Staged file storage 2) Database table storage 3) Fail-safe and Time Travel storage To use Snowsight to explore storage costs: 1. Navigate to Admin » Usage. 2. Select a warehouse to use to view the usage data. Snowflake recommends using an XS warehouse for this purpose. 3. Select Storage from the Usage Type drop-down. https://docs.snowflake.com/en/user-guide/cost-exploring-data-storagelatest Snowflake UI: Admin Usage Storage ( All Type Storage drop down list)

upvoted 1 times

**Marge23** 3 years, 1 month ago

latest Snowflake UI: Admin Usage Storage ( All Type Storage drop down list)

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**VisionBI** 3 years, 7 months ago

It should be Admin > Usage rather than Account > Usage & Billing

upvoted 1 times

**EmiB** 3 years, 8 months ago

Also: SNOWFLAKE.ACCOUNT_USAGE.STORAGE_USAGE

upvoted 2 times

**Tapasgup007** 3 years, 11 months ago

B is correct

upvoted 2 times

**Percy2112** 4 years, 2 months ago

B is correct Answer.

upvoted 1 times

**arpit_dataguy** 4 years, 9 months ago

**Selected Answer: B**

The correct answer is both B&D.

upvoted 3 times

**arpit_dataguy** 4 years, 9 months ago

The correct answer is both B&D.Sorry, only B is the correct answer.

upvoted 3 times

**arpit_dataguy** 4 years, 9 months ago

Sorry, only B is the correct answer.

upvoted 1 times

Question #9 *Topic 1*

Credit Consumption by the Compute Layer (Virtual Warehouses) is based on: (Choose two.)

A. Number of users

B. Warehouse size

C. Amount of data processed

D. # of clusters for the Warehouse

**Suggested Answer:** *B, D*

*Community vote distribution*

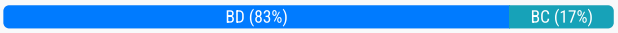


**jjordan** Highly Voted 4 years, 10 months ago

The answer to this is wrong. Consumption is based on the Size of the Warehouse, Number of Clusters and the amount of time the warehouse is running. The correct answer is therefore B&DI tend to disagree. Option D doesn't say active clusters and if the clusters are inactive they don't contribute to the consumption.

upvoted 27 times

**jjordan** Highly Voted 4 years, 10 months ago

The answer to this is wrong. Consumption is based on the Size of the Warehouse, Number of Clusters and the amount of time the warehouse is running. The correct answer is therefore B&D

upvoted 27 times

**blazingliverbird** 4 years, 9 months ago

I tend to disagree. Option D doesn't say active clusters and if the clusters are inactive they don't contribute to the consumption.

upvoted 7 times

**Wosama** 1 year, 4 months ago

B D IS CORRECT

upvoted 1 times

**jeny12138** 1 year, 7 months ago

BD is correct

upvoted 1 times

**SuryaG2490** 1 year, 8 months ago

B. Warehouse size: The size of the virtual warehouse directly impacts the number of credits consumed. Larger warehouses consume more credits per unit of time. D. # of clusters for the Warehouse: The number of clusters in a multi-cluster warehouse also affects credit consumption. More clusters mean higher credit usage.

upvoted 1 times

**datarrr** 1 year, 10 months ago

Selected Answer: BD

A and C are incorrect, as compute cost is based on wh size and the number or wh running

upvoted 1 times

**datarrr** 1 year, 10 months ago

A and C are incorrect, as compute cost is based on wh size and the number or wh runningB,D C is wrong because You don't pay for the "amount of data processed" but for how many TIME You spent for process the data e how cost by minute (that time).

upvoted 1 times

**BigDataBB** Most Recent 4 months, 1 week ago

B,D C is wrong because You don't pay for the "amount of data processed" but for how many TIME You spent for process the data e how cost by minute (that time).

upvoted 1 times

**junpal** 2 years, 1 month ago

**Selected Answer: BC**

B. Warehouse size: The size of the virtual warehouse is one of the primary factors that determine the credit consumption. Larger warehouses have more compute resources available, which allows them to process data faster and more efficiently, but also results in higher credit consumption.

C. Amount of data processed: The amount of data processed by the virtual warehouse is another key factor that determines credit consumption. Snowflake charges credits based on the amount of data processed, regardless of the number of users or the number of clusters used by the virtual warehouse.

A. Number of users and D. # of clusters for the warehouse are not direct factors in credit consumption by the compute layer. However, the number of users and the number of clusters used can indirectly affect credit consumption by impacting the amount of data processed and the warehouse size required to process that data efficiently.

Therefore, the credit consumption by the compute layer in Snowflake is primarily based on warehouse size and the amount of data processed.

upvoted 2 times

**nsyilmazer** 2 years ago

not sure why everything thinks it's about the number of clusters. Sizing determines the number of clusters and/or scaling option as well.

upvoted 1 times

**junpal** 2 years, 1 month ago

B. Warehouse size: The size of the virtual warehouse is one of the primary factors that determine the credit consumption. Larger warehouses have more compute resources available, which allows them to process data faster and more efficiently, but also results in higher credit consumption. C. Amount of data processed: The amount of data processed by the virtual warehouse is another key factor that determines credit consumption. Snowflake charges credits based on the amount of data processed, regardless of the number of users or the number of clusters used by the virtual warehouse. A. Number of users and D. # of clusters for the warehouse are not direct factors in credit consumption by the compute layer. However, the number of users and the number of clusters used can indirectly affect credit consumption by impacting the amount of data processed and the warehouse size required to process that data efficiently. Therefore, the credit consumption by the compute layer in Snowflake is primarily based on warehouse size and the amount of data processed.Here the question talks about the Compute Layer, you can't choose Answer "C". Correct answer are B and D.

upvoted 2 times

**location75** 3 years, 7 months ago

Here the question talks about the Compute Layer, you can't choose Answer "C". Correct answer are B and D.

upvoted 1 times

**PrashantGupta1616** 2 years, 1 month ago

The correct options for Snowflake Credit Consumption by the Compute Layer (Virtual Warehouses) are: B. Warehouse size D. # of clusters for the Warehouse Snowflake's credit consumption is influenced by the size of the virtual warehouse (which determines the amount of compute resources allocated) and the number of clusters used by the warehouse (especially in multi-cluster warehouses, which allow for scaling out to handle varying workloads).

upvoted 1 times

**damaldon** 2 years, 3 months ago

Ans: BD Credit charges are calculated based on: The warehouse size. The number of clusters (if using multi-cluster warehouses). The length of time the compute resources in each cluster runs.

upvoted 1 times

**ftizie** 2 years, 3 months ago

Amount of data has nothing to do with credit consumption; so the right answer is B and D, i.e. WH Size plus Number of clusters

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

It depends on warehouse size and no.of clusters for the warehouse

upvoted 1 times

**pranalig** 2 years, 4 months ago

correct answer:BD

upvoted 2 times

**Touchstone98** 2 years, 6 months ago

BD IS CORRECT

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

B &C is correct answer

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A.Warehouse size B.# of clusters for the Warehouse

upvoted 1 times

**whiteomax** 2 years, 9 months ago

After some research, I feel like picking C over D because you can have 8 2-node clusters or 4 4-node clusters in the same 16-node size warehouse. So more clusters should not cost you more. I also don't like C but generally more data means more time or resources used.

upvoted 1 times

**KHarish** 2 years, 11 months ago

B&D - warehouse size and no:of clusters and the amount of the time warehouse ran to process irrespective of the amount of the data processed

upvoted 1 times

Question #10 *Topic 1*

Which statement best describes `clustering`?

A. Clustering represents the way data is grouped together and stored within Snowflake's micro-partitions

B. The database administrator must define the clustering methodology for each Snowflake table

C. The clustering key must be included on the COPY command when loading data into Snowflake

D. Clustering can be disabled within a Snowflake account

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**jeny12138** Most Recent 1 year, 7 months ago

A is correct

upvoted 1 times

**tharangaweer** 1 year, 9 months ago

A.clustering represents the way data is grouped together and stored within sf`s micro partitions

upvoted 2 times

**Waseem111** 1 year, 10 months ago

Clustering represents the way data is grouped

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A. Clustering represents the way data is grouped together and stored within Snowflake's micro-partitions

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

gthyjtytyy

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

A.Clustering represents the way data is grouped together and stored within Snowflake's micro-partitions

upvoted 2 times

**_yyukta** 2 years, 8 months ago

A.Clustering represents the way data is grouped together and stored within Snowflake's micro-partitions

upvoted 1 times

**KHarish** 2 years, 11 months ago

A.Clustering represents the way data is grouped together and stored within Snowflake's micro-partitions

upvoted 1 times

**VS3497** 3 years, 4 months ago

A.Clustering represents the way data is grouped together and stored within Snowflake's micro-partitions

upvoted 1 times

**instride** 3 years, 5 months ago

A is correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

A is correct

upvoted 2 times

**bpxgeek** 4 years, 8 months ago

A IS TH GOOD ANSWER

upvoted 2 times

**arpit_dataguy** 4 years, 9 months ago

Strongly agree, A is the correct the correct answer.

upvoted 1 times

**Sundarr** 4 years, 9 months ago

https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions.html In Snowflake, as data is inserted/loaded into a table, clustering metadata is collected and recorded for each micro-partition created during the process. Snowflake then leverages this clustering information to avoid unnecessary scanning of micro-partitions during querying, significantly accelerating the performance of queries that reference these columns.

upvoted 2 times

Question #11 *Topic 1*

True or False: The COPY command must specify a File Format in order to execute.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

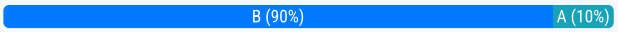


**jjordan** Highly Voted 4 years, 10 months ago

The correct answer is FALSE. The file format can be defined at the stage, table, or copy into command. Create Stage: https://docs.snowflake.com/en/sql-reference/sql/create-stage.html Create Table (STAGE_FILE_FORMAT option): https://docs.snowflake.com/en/sql-reference/sql/create-table.html Copy Into: https://docs.snowflake.com/en/sql-reference/sql/copy-into-table.html

upvoted 20 times

**tbridges** Highly Voted 3 years, 5 months ago

False. In Snowflake, the COPY command does not necessarily require specifying a File Format in order to execute. The COPY command can be used to load data into a table from various file formats such as CSV, JSON, Avro, Parquet, etc. When a file format is not explicitly specified in the COPY command, Snowflake attempts to automatically infer the file format based on the file's extension or internal file structure. This automatic file format inference feature is known as "automatic file format detection." However, if the file format cannot be determined automatically or if you want to override the automatic detection, you can explicitly specify a File Format in the COPY command using the FILE_FORMAT parameter. So, while specifying a File Format can provide more control and customization over the data loading process, it is not mandatory for the COPY command to execute successfully.

upvoted 8 times

**Shruthi34** Most Recent 1 year, 4 months ago

Answer is B. We can define file format while creating stage

upvoted 1 times

**d3697ab** 1 year, 5 months ago

The correct answer is B

upvoted 1 times

**Prakas1954** 1 year, 7 months ago

File format is not necessary

upvoted 1 times

**jeny12138** 1 year, 7 months ago

FALSE FALSE

upvoted 1 times

**cappper05** 1 year, 7 months ago

No need to specify a File

upvoted 1 times

**MikasSi** 1 year, 8 months ago

It is not mandatory to specify a FILE FORMAT for COPY operations. While it is not mandatory a FILE FORMAT should always be created for COPY into operations and file unloading and COPY into operations. Automatic inference should only be used with small sets of data and for testing or low importance data.

upvoted 1 times

**jiriz** 2 years, 1 month ago

False, it's possible but not neccesary.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A.True

upvoted 1 times

**pranalig** 2 years, 4 months ago

correct answer:B

upvoted 1 times

**JasMozai** 2 years, 6 months ago

It is optional

upvoted 1 times

**Stevec80** 2 years, 6 months ago

In the COPY command, the File_Format is not mandatory. With the Internal Named Stage you can specify file_format during creation of the stage

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

True-A

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A.True

upvoted 1 times

**KHarish** 2 years, 11 months ago

B is correct

upvoted 1 times

**annasri** 3 years ago

file format is nor mandatory

upvoted 1 times

Question #12 *Topic 1*

Which of the following commands sets the Virtual Warehouse for a session?

A. COPY WAREHOUSE FROM >;

B. SET WAREHOUSE = >;

C. USE WAREHOUSE >;

D. USE VIRTUAL_WAREHOUSE >;

**Suggested Answer:** *C*

*Community vote distribution*




**Deepika_s** Most Recent 1 year, 6 months ago

Although, USE WAREHOUSE and SET WAREHOUSE commands serve the same purpose -w.r.t. SNOWFLAKE we use USE WAREHOUSE command, so Ans is C.

upvoted 1 times

**jeny12138** 1 year, 7 months ago

c is correct

upvoted 1 times

**jeny12138** 1 year, 7 months ago

D or C

upvoted 1 times

**reachbharathan** 1 year, 7 months ago

Both USE WAREHOUSE and SET WAREHOUSE commands serve the same purpose - they switch the current session to use the specified warehouse. They are completely interchangeable and functionally identical. The only difference is syntactic preference, similar to how both USE SCHEMA and SET SCHEMA are equivalent commands. Example usage: USE WAREHOUSE compute_wh SET WAREHOUSE compute_wh Choose whichever syntax you prefer - they will produce the same result.

upvoted 1 times

**tharangaweer** 1 year, 9 months ago

c.use warehouse

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

C. USE WAREHOUSE >;

upvoted 1 times

**Vivek121** 2 years, 6 months ago

Why NOT B?

upvoted 2 times

**AnireddySaikiranReddy** 2 years, 8 months ago

C. USE WAREHOUSE >;

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C.USE WAREHOUSE > . C is correct

upvoted 1 times

**KHarish** 2 years, 11 months ago

USE WAREHOUSE > . C is correct

upvoted 1 times

**VS3497** 3 years, 4 months ago

C. USE WAREHOUSE >;

upvoted 1 times

**instride** 3 years, 5 months ago

The correct answer is C. According to snowflake docs, "A Snowflake session can only have one current warehouse at a time. The current warehouse for a session can be specified or changed at any time through the USE WAREHOUSE command."

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

C is correct

upvoted 1 times

**rajivgupta780** 3 years, 9 months ago

C is correct

upvoted 2 times

**Tapasgup007** 3 years, 11 months ago

C is correct

upvoted 1 times

**Percy2112** 4 years, 2 months ago

Use Warehouse command.

upvoted 1 times

**LuciBhish** 4 years, 8 months ago

USE WAREHOUSE ;

upvoted 3 times

Question #13 *Topic 1*

Which of the following objects can be cloned? (Choose four.)

A. Tables

B. Named File Formats

C. Schemas

D. Shares

E. Databases

F. Users

**Suggested Answer:** *A, B, C, E*

*Community vote distribution*

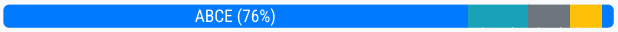


**AI_323** Highly Voted 3 years, 8 months ago

ABCE : Shares and Users cannot be cloned in Snowflake because they represent different types of objects and have different purposes. Shares represent shared access to data within Snowflake, and they cannot be cloned because access to data is controlled through the sharing mechanism, not through copying the share itself. Users represent individual accounts within Snowflake, and they cannot be cloned because each user account must have a unique name and must be associated with a unique set of permissions. Cloning a user would result in two users with the same name, which is not allowed in Snowflake. Instead of cloning Shares or Users, you can use the Snowflake GRANT command to manage access to data and permissions in Snowflake. The GRANT command allows you to control access to data by granting specific privileges to users, roles, or groups.+agree.

upvoted 24 times

**AI_323** Highly Voted 3 years, 8 months ago

ABCE : Shares and Users cannot be cloned in Snowflake because they represent different types of objects and have different purposes.

Shares represent shared access to data within Snowflake, and they cannot be cloned because access to data is controlled through the sharing mechanism, not through copying the share itself.

Users represent individual accounts within Snowflake, and they cannot be cloned because each user account must have a unique name and must be associated with a unique set of permissions. Cloning a user would result in two users with the same name, which is not allowed in Snowflake.

Instead of cloning Shares or Users, you can use the Snowflake GRANT command to manage access to data and permissions in Snowflake. The GRANT command allows you to control access to data by granting specific privileges to users, roles, or groups.

upvoted 24 times

**VLMBL_94** 2 years, 3 months ago

+agree.

upvoted 2 times

**srj_mehta** 3 years, 4 months ago

answer: ABCE Below is the list of all objects which can be CLONED** in Snowflake. Data Containment Objects Databases Schemas Tables Streams Data Configuration and Transformation Objects Stages File Formats Sequences Tasks Note: Internal named stages cannot be cloned.

upvoted 7 times

**jeny12138** Most Recent 1 year, 7 months ago

ABCE is correct

upvoted 1 times

**jeny12138** 1 year, 7 months ago

ABDF is correct

upvoted 1 times

**tharangaweer** 1 year, 9 months ago

Table,schemas,file formats, databases can be cloned

upvoted 1 times

**shubhat** 1 year, 9 months ago

Complete list is as follows: ○ Databases ○ Schemas ○ Tables ○ Views (including materialized views) ○ Streams ○ Stages ○ Pipes ○ Tasks ○ File Formats ○ Sequences

upvoted 1 times

**aba2s** 2 years ago

D is False because Shares cannot be cloned or reshared F is also false

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A,C,E ARE CORRECT ANSWER

upvoted 1 times

**pranalig** 2 years, 4 months ago

ABCE: Users and schema cannot be clone

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

The following account-level objects cannot be cloned: Users Roles Grants Virtual Warehouses Resource monitors Storage integrations

upvoted 1 times

**JasMozai** 2 years, 6 months ago

Creates a copy of an existing object in the system. This command is primarily used for creating zero-copy clones of databases, schemas, and tables; however, it can also be used to quickly/easily create clones of other schema objects , such as external stages, file formats, and sequences, and database roles.

upvoted 1 times

**Stevec80** 2 years, 6 months ago

https://docs.snowflake.com/en/sql-reference/sql/create-clone

upvoted 1 times

**Caesarpanda** 2 years, 7 months ago

Below are the items we can clone according to chatgpt : databases schemas tables views MV Shares Roles Stored Proc User Defined Func Seq Streams So the correct answer is A,C,D,E

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

ACEF are correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A B C E

upvoted 1 times

**VenkataPhaniPavanKumar** 2 years, 9 months ago

ABCE is the right answer To break it down Database, Schemas and Tables can be clone(no brainer). Users cannot be cloned like the above objects. Shares cannot be clone and no where in the docs mentioned about this cloning. The only one remaining from the choice is Named File Format.

upvoted 1 times

**Kritig08** 2 years, 10 months ago

User can't be cloned like database objects. As best practice -- this is from snowflake community.

upvoted 1 times

Question #14 *Topic 1*

Which object allows you to limit the number of credits consumed within a Snowflake account?

A. Account Usage Tracking

B. Resource Monitor

C. Warehouse Limit Parameter

D. Credit Consumption Tracker

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Percy2112** Highly Voted 4 years, 2 months ago

We can control and set the limit for credit usage in Resource Monitors. https://docs.snowflake.com/en/user-guide/resource-monitors.html#:~:text=To%20help%20control%20costs%20and%20avoid%20unexpected%20credit%20usage%20caused%20by%20running%20warehouses%2C%20Snowflake%20provides%20resource%20monitors

upvoted 7 times

**jeny12138** Most Recent 1 year, 7 months ago

B is correct

upvoted 1 times

**aba2s** 2 years ago

A resource monitor can help control costs and avoid unexpected credit usage caused by running warehouses. A virtual warehouse consumes Snowflake credits while it runs. You can use a resource monitor to monitor credit usage by virtual warehouses and the cloud services needed to support those warehouses. You can also set up a resource monitor to suspend a user-managed virtual warehouse when it reaches a credit limit.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B. Resource Monitor

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

B is correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Resource Monitor

upvoted 1 times

**SamuelZhang** 3 years, 3 months ago

B is correct

upvoted 1 times

**VS3497** 3 years, 4 months ago

D. Credit Consumption Tracker

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

B is correct

upvoted 2 times

**grabc** 4 years, 9 months ago

B is correctTrue B is correct

upvoted 2 times

**grabc** 4 years, 9 months ago

B is correct

upvoted 2 times

**bpxgeek** 4 years, 8 months ago

True B is correct

upvoted 1 times

Question #15 *Topic 1*

Snowflake is designed for which type of workloads? (Choose two.)

A. OLAP (Analytics) workloads

B. OLTP (Transactional) workloads

C. Concurrent workloads

D. On-premise workloads

**Suggested Answer:** *A, C*

*Community vote distribution*

AC (84%) AB (13%)

**Bigdata4youall** Highly Voted 4 years, 10 months ago

I think the answer is A and C, OLTP is not best use case for snowflake https://www.quora.com/Can-Snowflake-be-used-for-an-OLTP-system-or-is-it-only-best-suited-for-warehousing

upvoted 10 times

**jeny12138** Most Recent 1 year, 7 months ago

noncurrent and analytics

upvoted 1 times

**jeny12138** 1 year, 7 months ago

B?C is corrct

upvoted 1 times

**Yash_Chavan** 1 year, 9 months ago

Snowflake is designed for OLAP (Analytics) workloads and Concurrent workloads.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

OLAP AND Concurrent Workloads

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

A. OLAP (Analytics) workloads C. Concurrent workloads Snowflake is primarily designed for analytical processing (OLAP) and can efficiently handle concurrent workloads due to its architecture, which allows multiple virtual warehouses to operate independently and scale as needed. It is not optimized for OLTP (Transactional) workloads, which typically require low-latency operations, nor is it designed specifically for on-premise workloads, as it is a cloud-based data warehousing solution.

upvoted 3 times

**AnireddySaikiranReddy** 2 years, 8 months ago

A.OLAP (Analytics) workloads B.Concurrent workloads

upvoted 2 times

**_yyukta** 2 years, 8 months ago

A.OLAP (Analytics) workloads B.Concurrent workloads

upvoted 1 times

**mjspnieo** 2 years, 11 months ago

A and B Snowflake Unistore offers modern approach to working with transactional and analytical data together in a single platform. https://www.snowflake.com/guides/olap-vs-oltp

upvoted 2 times

**mjspnieo** 2 years, 11 months ago

B and C Snowflake Unistore offers modern approach to working with transactional and analytical data together in a single platform https://www.snowflake.com/guides/olap-vs-oltp

upvoted 1 times

**ifmateus** 3 years, 2 months ago

The correct answer is OLAP and OLTP. Why? Because Snf introduced a new feature "Unistore" which enables hybrid tables with both transactional type data (OLTP) and analytical data (OLAP), in the PAST the answer would be different as one can see in the answer given by Snowflake below. https://community.snowflake.com/s/question/0D50Z00008BDIPTSA5/snowflake-oltp https://www.snowflake.com/en/data-cloud/workloads/unistore/

upvoted 2 times

**VS3497** 3 years, 4 months ago

A. OLAP (Analytics) workloads B. OLTP (Transactional) workloads

upvoted 2 times

**danlsn** 3 years, 4 months ago

As a Data Warehouse it's surely optimised for OLAP and concurrent workloads.

upvoted 1 times

**Nitin001** 3 years, 5 months ago

Check this new feature of Unistore! A new workload that delivers a modern approach to working with transactional and analytical data together in a single platform https://www.snowflake.com/blog/introducing-unistore/

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

AB is correct

upvoted 1 times

**piconcomeback** 3 years, 6 months ago

Part of the Data Cloud, Snowflake Unistore offers modern approach to working with transactional and analytical data together in a single platform. https://www.snowflake.com/guides/olap-vs-oltp

upvoted 1 times

**jonathanb21** 3 years, 6 months ago

AC is the correct answer. - https://www.snowflake.com/trending/online-analytical-processing - https://www.snowflake.com/guides/olap-vs-oltp

upvoted 2 times

Question #16 *Topic 1*

What are the three layers that make up Snowflake's architecture? (Choose three.)

A. Compute

B. Tri-Secret Secure

C. Storage

D. Cloud Services

**Suggested Answer:** *A, C, D*

*Community vote distribution*

ACD (100%)

**Prakas1954** Most Recent 1 year, 7 months ago

Cloud, Compute, storage (CCS)

upvoted 2 times

**jeny12138** 1 year, 7 months ago

ACD is correct

upvoted 2 times

**tharangaweer** 1 year, 9 months ago

Storage,compute,cloud

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

Cloud Service,Compute and Storage Layer

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

A. Compute C. Storage D. Cloud Services

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. C. D.

upvoted 1 times

**VS3497** 3 years, 4 months ago

A. Compute C. Storage D. Cloud Services

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

ACD is the correct answer

upvoted 1 times

**EmiB** 3 years, 8 months ago

B - relates to security

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

ACD is correct

upvoted 1 times

**Percy2112** 4 years, 2 months ago

Easy and simple. We can't make mistake here.

upvoted 1 times

**harsh29** 4 years, 7 months ago

Snowflake architecture is made up of compute , storage and cloud services layer primarily.

upvoted 1 times

**LuciBhish** 4 years, 8 months ago

Storage, Compute and Cloud services

upvoted 3 times

**grabc** 4 years, 9 months ago

should be A,C,D

upvoted 1 times

Question #17 *Topic 1*

Why would a customer size a Virtual Warehouse from an X-Small to a Medium?

A. To accommodate more queries

B. To accommodate more users

C. To accommodate fluctuations in workload

D. To accommodate a more complex workload

**Suggested Answer:** *D*

*Community vote distribution*

D (96%) (4

**ggww** Highly Voted 4 years, 10 months ago

should be D

upvoted 20 times

**SteveKarr** Highly Voted 4 years, 8 months ago

A is the correct answer.

Ref: https://docs.snowflake.com/en/user-guide/warehouses-considerations.html

Resizing a warehouse generally improves query performance, particularly for larger, more complex queries. It can also help reduce the queuing that occurs if a warehouse does not have enough compute resources to process all the queries that are submitted concurrently. Note that warehouse resizing is not intended for handling concurrency issues; instead, use additional warehouses to handle the workload or use a multi-cluster warehouse (if this feature is available for your account).

upvoted 15 times

**SteveKarr** 4 years, 8 months ago

A is the correct answer. Ref: https://docs.snowflake.com/en/user-guide/warehouses-considerations.html Resizing a warehouse generally improves query performance, particularly for larger, more complex queries. It can also help reduce the queuing that occurs if a warehouse does not have enough compute resources to process all the queries that are submitted concurrently. Note that warehouse resizing is not intended for handling concurrency issues; instead, use additional warehouses to handle the workload or use a multi-cluster warehouse (if this feature is available for your account).You answer option contradicts you post. What you posted mentioned that "...particularly for larger, more complex queries.". So correct answer is DThe text you quote explains why D is the correct answer.The answer literally justifies why it's A.. Scaling up improves query performance and can reduce queing to process queries submitted concurrently. "Use additional warehouses to handle the workload". Everything related to workload is solved through Scaling Out (more VW) and not by increasing the size of the VW.Answer is D. As you mention warehouse resizing is not intended for handling concurrency issues(i.e handling more queries). The solution for A would be multi-clustering.

upvoted 15 times

**asorock** 3 years, 3 months ago

You answer option contradicts you post. What you posted mentioned that "...particularly for larger, more complex queries.". So correct answer is D

upvoted 8 times

**YokoSumiGaeshi** 3 years, 9 months ago

The text you quote explains why D is the correct answer.

upvoted 10 times

**ifmateus** 3 years, 2 months ago

The answer literally justifies why it's A.. Scaling up improves query performance and can reduce queing to process queries submitted concurrently. "Use additional warehouses to handle the workload". Everything related to workload is solved through Scaling Out (more VW) and not by increasing the size of the VW.

upvoted 2 times

**AdolinKholin** 3 years, 4 months ago

Answer is D. As you mention warehouse resizing is not intended for handling concurrency issues(i.e handling more queries). The solution for A would be multi-clustering.

upvoted 5 times

**jeny12138** Most Recent 1 year, 7 months ago

D is correct

upvoted 1 times

**tharangaweer** 1 year, 9 months ago

To accomodate more complex queries

upvoted 1 times

**mackel** 2 years, 1 month ago

should be D: "Note that warehouse resizing is not intended for handling concurrency issues"; i.e. number of queries...

upvoted 1 times

**VLMBL_94** 2 years, 3 months ago

Acc to snowflake document "In general, you should try to match the size of the warehouse to the expected size and complexity of the queries to be processed by the warehouse"

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

To Accommodate a More workloads

upvoted 1 times

**nsvsai1357** 2 years, 3 months ago

D is correct answer If queries processed by a warehouse are running slowly, you can always resize the warehouse to provision more compute resources. The additional resources do not impact any queries that are already running, but once they are fully provisioned they become available for use by any queries that are queued or newly submitted.

upvoted 1 times

**pranalig** 2 years, 4 months ago

correct answer:D

upvoted 1 times

**p22_nilesh** 2 years, 5 months ago

Scale up is for complex queries

upvoted 1 times

**JasMozai** 2 years, 6 months ago

Everything related to workload is solved through Scaling Out (more VW) and not by increasing the size of the VW.

upvoted 1 times

**Chhatresh** 2 years, 7 months ago

Resizing a warehouse to a larger size is useful when the operations being performed by the warehouse will benefit from more compute resources, including: Improving the performance of large, complex queries against large data sets. Improving performance while loading and unloading significant amounts of data.

upvoted 2 times

**AnireddySaikiranReddy** 2 years, 8 months ago

The answer id D

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D is correct

upvoted 1 times

**e_l_i_a_s** 2 years, 9 months ago

The answer id D

upvoted 1 times

**avashish** 2 years, 10 months ago

D is correct answer because when query is more complex then it will be using more compute power so in that case VW size should be increased from XS to medium

upvoted 1 times

**SolutionA** 3 years ago

What is the definition of complex workload the option is incomplete

upvoted 1 times

Question #18 *Topic 1*

True or False: Reader Accounts incur no additional Compute costs to the Data Provider since they are simply reading the shared data without making changes.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*




**jeny12138** Most Recent 1 year, 7 months ago

it is false

upvoted 2 times

**tharangaweer** 1 year, 9 months ago

The readers accounts usage cost applies to data provider

upvoted 3 times

**Mallikharjuna452** 2 years, 3 months ago

B.False

upvoted 1 times

**pranalig** 2 years, 4 months ago

correct answer:B

upvoted 1 times

**gunasekarantjk1996** 2 years, 5 months ago

False. While Reader Accounts in Snowflake do allow consumers to access and query shared data, the Data Provider assumes all responsibility for credit charges incurred by users in the Reader Account. This includes compute costs associated with executing queries12. So, even though the consumers are just reading the shared data, the compute resources they use for this purpose do result in additional costs for the Data Provider

upvoted 2 times

**AnireddySaikiranReddy** 2 years, 8 months ago

B. Note

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. False

upvoted 1 times

**0e504b5** 2 years, 9 months ago

https://docs.snowflake.com/en/user-guide/data-sharing-reader-create The reader account is created, owned, and managed by the provider account, which assumes all responsibility for credit charges incurred by users in the reader account. Similar to standard consumer accounts, the provider account uses shares to share databases with reader accounts; however, a reader account can only consume data from the provider account that created it. Warehouses in a reader account can consume an unlimited number of credits each month, which will be charged to your provider account. To limit usage, set up a resource monitor for the warehouse.

upvoted 3 times

**nonamae** 3 years, 1 month ago

The question is if the reader account does not incur additional charges to the data provider, which is false. Based from this statement "The reader account is created, owned, and managed by the provider account, which assumes all responsibility for credit charges incurred by users in the reader account." the provider account is responsible for any charges incurred by user in reader account, thus, there can be charges (which can be done by a user in a reader account) and the amount will be incurred on the one who created the reader account / the data provider

upvoted 2 times

**Marge23** 3 years, 1 month ago

B. Note Warehouses in a reader account can consume an unlimited number of credits each month, which will be charged to your provider account. To limit usage, set up a resource monitor for the warehouse.

upvoted 1 times

**Amitsnowflake** 3 years, 2 months ago

Reader accounts are outside Snowflake and cannot be charged.You are right , so data provider should pay! So once the reader account run a query it incurs additional cost for data provider!

upvoted 1 times

**Amitsnowflake** 3 years, 2 months ago

**Selected Answer: A**

Reader accounts are outside Snowflake and cannot be charged.

upvoted 1 times

**BobFar** 3 years ago

You are right , so data provider should pay! So once the reader account run a query it incurs additional cost for data provider!

upvoted 1 times

**Hariumesh** 3 years, 3 months ago

the answer is True According to "The reader account is created, owned, and managed by the provider account, which assumes all responsibility for credit charges incurred by users in the reader account. Similar to standard consumer accounts, the provider account uses shares to share databases with reader accounts; however, a reader account can only consume data from the provider account that created it: "

upvoted 2 times

**asorock** 3 years, 3 months ago

B is the Answer Warehouses in a reader account can consume an unlimited number of credits each month, which will be charged to your provider account. To limit the usage, please setup a resource monitor for the warehouse.

https://docs.snowflake.com/en/user-guide/data-sharing-reader-create

upvoted 2 times

**asorock** 3 years, 3 months ago

B is the Answer Warehouses in a reader account can consume an unlimited number of credits each month, which will be charged to your provider account. To limit the usage, please setup a resource monitor for the warehouse. https://docs.snowflake.com/en/user-guide/data-sharing-reader-createI meant to write A

upvoted 2 times

**asorock** 3 years, 3 months ago

I meant to write A

upvoted 1 times

**PoornaChandraIn** 3 years, 4 months ago

Answer is B, if you reading the data from reader account no charges applies here but if you are sharing the data what you read then charges applies.

upvoted 3 times

**Def21** 3 years, 4 months ago

B "The reader account is created, owned, and managed by the provider account, which assumes all responsibility for credit charges incurred by users in the reader account." https://docs.snowflake.com/en/user-guide/data-sharing-reader-create

upvoted 2 times

**VS3497** 3 years, 4 months ago

A. True

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 2 times

Question #19 *Topic 1*

Which of the following connectors allow Multi-Factor Authentication (MFA) authorization when connecting? (Choose all that apply.)

A. JDBC

B. SnowSQL

C. Snowflake Web Interface (UI)

D. ODBC

E. Python

**Suggested Answer:** *A, B, C, D, E*

*Community vote distribution*




**EmiB** Highly Voted 3 years, 8 months ago

MFA login is designed primarily for connecting to Snowflake through the - web interface - but is also fully-supported by SnowSQL - and the Snowflake JDBC and ODBC drivers. Snowflake supports MFA token caching with the following drivers and connectors on macOS and Windows. This feature is not supported on Linux. - ODBC driver version 2.23.0 (or later). - JDBC driver version 3.12.16 (or later). - Python Connector for Snowflake version 2.3.7 (or later).

upvoted 7 times

**alexhparreira** Most Recent 5 months ago

this question is tricky, because ODBC and JDBC are not connectors, but drivers

upvoted 1 times

**gadhepatil7** 10 months ago

https://docs.snowflake.com/en/user-guide/security-mfa#connecting-to-snowflake-with-mfa first line in this link doc all are supported

upvoted 1 times

**KaranSnow** 11 months ago

All are valid answers.

upvoted 1 times

**Prakas1954** 1 year, 7 months ago

MFA can be used with various Snowflake interfaces, including the web interface, SnowSQL, JDBC, Node.js, and ODBC drivers

upvoted 1 times

**jeny12138** 1 year, 7 months ago

all are correct

upvoted 1 times

**cappper05** 1 year, 7 months ago

no Python on the list

upvoted 1 times

**Yash_Chavan** 1 year, 9 months ago

Connector word is used to divert. Ans is All of them- ABCDE

upvoted 1 times

**pia33** 2 years, 2 months ago

E - only python is a connector

upvoted 1 times

**ABRocks** 2 years, 2 months ago

Its ADE, these are the connectors.

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

All options are correct

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

All(A,B,C,d,e)

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A B C D E

upvoted 1 times

**suhovet** 2 years, 10 months ago

ADE - are connectors

upvoted 3 times

**PriyankaTN** 3 years, 4 months ago

This question came in my exam and asked to choose two option..There was not UI option but have Spark instead of UI..

upvoted 4 times

**Mayuri917233** 3 years, 5 months ago

ABCDE is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

**Selected Answer: ABCDE**

All are correct

upvoted 3 times

**RajuNaik** 3 years, 6 months ago

https://docs.snowflake.com/en/user-guide/security-mfa

upvoted 2 times

Question #20 *Topic 1*

True or False: Snowflake charges a premium for storing semi-structured data.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**jeny12138** Most Recent 1 year, 7 months ago

false. standard also works

upvoted 1 times

**tharangaweer** 1 year, 9 months ago

irrespective of structured or semi structured

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B.False,Snowflake does not charge a premium for storing semi-structured data

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

False is corect

upvoted 1 times

**_yyukta** 2 years, 8 months ago

b false is correct

upvoted 1 times

**PoornaChandraIn** 3 years, 4 months ago

Answer is B. Snowflake charges as per its standards irrespective of structured or semi structured data.

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 2 times

**abbcia** 3 years, 7 months ago

it's b

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

B is correct

upvoted 1 times

**Percy2112** 4 years, 2 months ago

Snowflake does not charge extra for any special/semi-structure type of file.

upvoted 2 times

**Sundarr** 4 years, 9 months ago

B is the answer

upvoted 3 times

Question #21 *Topic 1*

Which of the following statements describes a benefit of Snowflake's separation of compute and storage? (Choose all that apply.)

A. Growth of storage and compute are tightly coupled together

B. Storage expands without the requirement to add more compute

C. Compute can be scaled up or down without the requirement to add more storage

D. Multiple compute clusters can access stored data without contention

**Suggested Answer:** *B, C, D*

*Community vote distribution*

BCD (100%)

**Sly007** Highly Voted 4 years, 6 months ago

BCD for sure

upvoted 6 times

**jeny12138** Most Recent 1 year, 7 months ago

Snowflake's separation decouples compute and storage

upvoted 1 times

**tharangaweer** 1 year, 9 months ago

bcd are correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B,C,D ARE CORRECT OPTIONS

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

BCD are correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

BCD are correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

BCD are correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

BCD are correct

upvoted 3 times

**MultiCloudIronMan** 4 years ago

Is the right answer

upvoted 1 times

**Percy2112** 4 years, 2 months ago

BCD are the correct options.

upvoted 1 times

**Graghu** 4 years, 6 months ago

BCD are the right options

upvoted 2 times

**bpxgeek** 4 years, 8 months ago

BCD is the good answer

upvoted 2 times

**Jialu** 4 years, 8 months ago

Answer are BCD

upvoted 2 times

**Sid460545** 4 years, 9 months ago

BCD, all the three are correct

upvoted 3 times

**Sundarr** 4 years, 9 months ago

Answer are BCD

upvoted 2 times

**rsmurugan** 4 years, 9 months ago

Answer are BCD

upvoted 3 times

**Hello1991** 4 years, 9 months ago

Answer are BCD

upvoted 3 times

Question #22 *Topic 1*

True or False: It is possible to unload structured data to semi-structured formats such as JSON and Parquet.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**grabc** Highly Voted 4 years, 9 months ago

A is correct: https://docs.snowflake.com/en/user-guide/data-unload-considerations.html#unloading-a-relational-table-to-json

upvoted 7 times

**jeny12138** Most Recent 1 year, 7 months ago

True is correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A.True

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

A True

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A True

upvoted 1 times

**avashish** 2 years, 9 months ago

A is correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

A is correct

upvoted 1 times

**Percy2112** 4 years, 2 months ago

True. Correct Answer.

upvoted 1 times

**Diegovaz** 4 years, 4 months ago

CSV is also an option!

upvoted 1 times

**Surbin** 4 years, 6 months ago

A is the correct answer

upvoted 2 times

Question #23 *Topic 1*

In which layer of its architecture does Snowflake store its metadata statistics?

A. Storage Layer

B. Compute Layer

C. Database Layer

D. Cloud Services Layer

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**jeny12138** Most Recent 1 year, 7 months ago

cloud services layer is for meta data

upvoted 2 times

**M477h3w** 1 year, 7 months ago

metadata management is under "cloud services" https://docs.snowflake.com/en/user-guide/intro-key-concepts

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

D.Cloud Service Layer

upvoted 2 times

**subhataanu** 2 years, 5 months ago

D https://docs.snowflake.com/en/user-guide/intro-key-concepts#database-storage Snowflake manages all aspects of how this data is stored — the organization, file size, structure, compression, metadata, statistics, and other aspects of data storage are handled by Snowflake. The data objects stored by Snowflake are not directly visible nor accessible by customers; they are only accessible through SQL query operations run using Snowflake.

upvoted 1 times

**subhataanu** 2 years, 5 months ago

D https://docs.snowflake.com/en/user-guide/intro-key-concepts#database-storage Snowflake manages all aspects of how this data is stored — the organization, file size, structure, compression, metadata, statistics, and other aspects of data storage are handled by Snowflake. The data objects stored by Snowflake are not directly visible nor accessible by customers; they are only accessible through SQL query operations run using Snowflake.Typo - Answer is C

upvoted 1 times

**subhataanu** 2 years, 5 months ago

Typo - Answer is C

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

D is correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D is correct

upvoted 1 times

**avashish** 2 years, 9 months ago

D cloud Service Layer

upvoted 1 times

**zephyrvz** 3 years ago

D cloud service layer

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**harshkoshta** 3 years, 4 months ago

I think option A is correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

D Cloud Service Layer

upvoted 1 times

**location75** 3 years, 7 months ago

D Cloud Service Layer of cause

upvoted 1 times

**Snovice** 3 years, 7 months ago

Correct answer should be "C". https://docs.snowflake.com/en/user-guide/intro-key-concepts Database Storage¶ When data is loaded into Snowflake, Snowflake reorganizes that data into its internal optimized, compressed, columnar format. Snowflake stores this optimized data in cloud storage. Snowflake manages all aspects of how this data is stored — the organization, file size, structure, compression, metadata, statistics, and other aspects of data storage are handled by Snowflake. The data objects stored by Snowflake are not directly visible nor accessible by customers; they are only accessible through SQL query operations run using Snowflake.

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

D is correct

upvoted 1 times

**AB_DA** 4 years, 1 month ago

Cloud service layer

upvoted 1 times

**Percy2112** 4 years, 2 months ago

Cloud Service layer stores metadata of queries for optimization.

upvoted 1 times

**Sundarr** 4 years, 9 months ago

Metadata manager is part of Cloud Services layer.

upvoted 2 times

Question #24 *Topic 1*

True or False: Data in fail-safe can be deleted by a user or the Snowflake team before it expires.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**jeny12138** Most Recent 1 year, 7 months ago
cannot be delete
upvoted 1 times

**SergesA** 1 year, 9 months ago
no user action provide
upvoted 1 times

**aemilka** 2 years, 2 months ago
Fail-safe provides a (non-configurable) 7-day period during which historical data may be recoverable by Snowflake. No user operations allowed.
https://docs.snowflake.com/en/user-guide/data-failsafe
upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago
The User can't have access to Fail Safe,Snowflake Manage the Fail-safe
upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago
Answer is B as user doesn't have access/control to modify fail-safe.
upvoted 1 times

**_yyukta** 2 years, 8 months ago
b false
upvoted 1 times

**Avees** 3 years ago
user does not have control or access over fail-safe
upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago
Correct
upvoted 1 times

**PoornaChandraIn** 3 years, 4 months ago
Answer is B as user doesn't have access/control to modify fail-safe.
upvoted 1 times

**Mayuri917233** 3 years, 5 months ago
B is correct
upvoted 1 times

**Tapasgup007** 3 years, 11 months ago
B is correct
upvoted 1 times

**manohar3241** 4 years ago
Failsafe manages by Snowflake support only
upvoted 2 times

**Percy2112** 4 years, 2 months ago
Fail Safe can not be configured. Users, neither retrieve nor delete the data of fail safe.
upvoted 1 times

**learnazureportal** 4 years, 5 months ago

Answer is B. Failsafe manages by Snowflake support not user.

upvoted 1 times

**grabc** 4 years, 9 months ago

It doesn't make sense for user to remove date in fail-safe, B is correct

upvoted 2 times

Question #25 *Topic 1*

True or False: Snowflake's data warehouse was built from the ground up for the cloud in lieu of using an existing database or a platform, like Hadoop, as a base.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*




**sdnv89** Highly Voted 3 years, 11 months ago

'in lieu of' == 'instead of'. Yes, Snowflake was built from the ground up, without using something else as a base. So the statement is true and the answer is A.

upvoted 38 times

**jeny12138** Most Recent 1 year, 7 months ago

unsure about this

upvoted 1 times

**al.kepner** 2 years, 2 months ago

Why is the Correct answer claimed to be False?

upvoted 3 times

**pia33** 2 years, 2 months ago

A is correct

upvoted 1 times

**ABRocks** 2 years, 2 months ago

True, it's a mix of shared nothing and shared disk.

upvoted 2 times

**hpsk71** 2 years, 2 months ago

Option B- False- snowflake was not built in lieu of platforms like hadoop

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B.False

upvoted 1 times

**pranalig** 2 years, 4 months ago

correct answer:A

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

**True.** Snowflake's data warehouse was built from the ground up specifically for the cloud, rather than using an existing database or platform like Hadoop as a base. This design allows Snowflake to leverage the cloud's scalability, elasticity, and performance features, providing a modern data warehousing solution tailored for cloud environments.

upvoted 2 times

**sambathchandran** 2 years, 5 months ago

Snowflake's data warehouse was built from the ground up specifically for the cloud, rather than being based on an existing database or platform like Hadoop. This design allows Snowflake to fully leverage the scalability, elasticity, and performance benefits of cloud infrastructure, offering unique features such as automatic scaling, data sharing, and a multi-cluster architecture. A.true

upvoted 2 times

**AnireddySaikiranReddy** 2 years, 8 months ago

b False

upvoted 1 times

**_yyukta** 2 years, 8 months ago

b False

upvoted 1 times

**rekibjec10** 2 years, 9 months ago

B https://docs.snowflake.com/en/user-guide/intro-key-concepts

upvoted 1 times

**rekibjec10** 2 years, 9 months ago

B https://docs.snowflake.com/en/user-guide/intro-key-conceptsthat link says this in the second paragraph: "The Snowflake data platform is not built on any existing database technology or “big data” software platforms such as Hadoop. Instead, Snowflake combines a completely new SQL query engine with an innovative architecture natively designed for the cloud. "

upvoted 1 times

**56d43d6** 1 year, 10 months ago

that link says this in the second paragraph: "The Snowflake data platform is not built on any existing database technology or “big data” software platforms such as Hadoop. Instead, Snowflake combines a completely new SQL query engine with an innovative architecture natively designed for the cloud. "

upvoted 1 times

**sandy91** 2 years, 9 months ago

in lieu = instead of

upvoted 1 times

**Isio05** 2 years, 12 months ago

For sure A

upvoted 1 times

**zephyrvz** 3 years ago

A is correct

upvoted 2 times

**BobFar** 3 years ago

A is true. In lieu = instead of The Snowflake data platform is not built on any existing database technology or “big data” software platforms such as Hadoop. Instead, Snowflake combines a completely new SQL query engine with an innovative architecture natively designed for the cloud. https://docs.snowflake.com/en/user-guide/intro-key-concepts.html

upvoted 3 times

Question #26 *Topic 1*

Which of the following statements are true of Virtual Warehouses? (Choose all that apply.)

A. Customers can change the size of the Warehouse after creation

B. A Warehouse can be resized while running

C. A Warehouse can be configured to suspend after a period of inactivity

D. A Warehouse can be configured to auto-resume when new queries are submitted

**Suggested Answer:** *A, B*

*Community vote distribution*

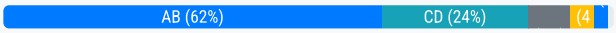


**ggww** Highly Voted 4 years, 10 months ago

all fours are truehttps://docs.snowflake.com/en/user-guide/warehouses-overviewCorrect, you can configure auto suspend and auto resume: https://docs.snowflake.com/en/user-guide/warehouses-considerations#automating-warehouse-suspension AND you can resize a wh even when it is running: https://docs.snowflake.com/en/user-guide/warehouses-overviewAdditionally, multi-cluster warehouses support all the same properties and actions as single-cluster warehouses, including: Specifying a warehouse size. Resizing a warehouse at any time. Auto-suspending a running warehouse due to inactivity; note that this does not apply to individual clusters, but rather the entire multi-cluster warehouse. Auto-resuming a suspended warehouse when new queries are submitted.

upvoted 140 times

**ggww** Highly Voted 4 years, 10 months ago

all fours are true

upvoted 138 times

**BobFar** 3 years ago

https://docs.snowflake.com/en/user-guide/warehouses-overview

upvoted 1 times

**56d43d6** 1 year, 10 months ago

Correct, you can configure auto suspend and auto resume: https://docs.snowflake.com/en/user-guide/warehouses-considerations#automating-warehouse-suspension AND you can resize a wh even when it is running: https://docs.snowflake.com/en/user-guide/warehouses-overview

upvoted 5 times

**Marge23** 3 years, 1 month ago

Additionally, multi-cluster warehouses support all the same properties and actions as single-cluster warehouses, including: Specifying a warehouse size. Resizing a warehouse at any time. Auto-suspending a running warehouse due to inactivity; note that this does not apply to individual clusters, but rather the entire multi-cluster warehouse. Auto-resuming a suspended warehouse when new queries are submitted.

upvoted 5 times

**Tapasgup007** 3 years, 11 months ago

All four are True

upvoted 24 times

**niho1993** Most Recent 1 year, 1 month ago

all fours are true

upvoted 2 times

**Zjrab** 1 year, 1 month ago

ABCD are true

upvoted 1 times

**Sk_3** 1 year, 1 month ago

Actually A,B,C,D all are correct

upvoted 1 times

**visionary4ever** 1 year, 2 months ago

All four are true A: resizing allowed B: Dynamic Scaling allowed C , D : Auto resume and auto suspends are also standard warehouse features

upvoted 1 times

**mishrask06** 1 year, 2 months ago

All four are true

upvoted 1 times

**f7dbe29** 1 year, 3 months ago

all are true

upvoted 1 times

**nizmo63** 1 year, 3 months ago

All are true

upvoted 1 times

**HagarTheHorrible** 1 year, 4 months ago

all true, not only two

upvoted 1 times

**slushmush** 1 year, 5 months ago

All four are correct, but only two are allowed to be selected by this question. That needs to be fixed.

upvoted 1 times

**Vijaya15** 1 year, 5 months ago

all fours are true

upvoted 1 times

**Prakas1954** 1 year, 7 months ago

all the answers are correct

upvoted 1 times

**jeny12138** 1 year, 7 months ago

four are all true

upvoted 1 times

**SuryaG2490** 1 year, 8 months ago

All of the following statements are true of Virtual Warehouses in Snowflake: A. Customers can change the size of the Warehouse after creation: You can resize a warehouse at any time to adjust the compute resources. B. A Warehouse can be resized while running: Warehouses can be resized even while they are running to accommodate changing workloads. C. A Warehouse can be configured to suspend after a period of inactivity: You can set a warehouse to automatically suspend after a specified period of inactivity to save costs. D. A Warehouse can be configured to auto-resume when new queries are submitted: Warehouses can be configured to automatically resume when new queries are submitted, ensuring that compute resources are available when needed.

upvoted 1 times

**MultiCloudIronMan** 1 year, 9 months ago

All are correct ABCD

upvoted 2 times

**datarrr** 1 year, 10 months ago

all fours are true, BUT the question only allows to pick up two options

upvoted 1 times

Question #27 *Topic 1*

The PUT command: (Choose two.)

A. Automatically creates a File Format object

B. Automatically uses the last Stage created

C. Automatically compresses files using Gzip

D. Automatically encrypts files

**Suggested Answer:** *C, D*

*Community vote distribution*

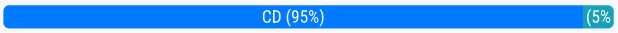


**sabithaMR** Most Recent 1 year, 6 months ago

CD are true

upvoted 2 times

**jeny12138** 1 year, 7 months ago

CD is true

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

The PUT command in Snowflake automatically compresses files using Gzip and encrypts the files during the transfer process

upvoted 3 times

**pranalig** 2 years, 4 months ago

correct answer:CD

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

C. Automatically compresses files using Gzip D. Automatically encrypts files Therefore, the correct options are: C. Automatically compresses files using Gzip D. Automatically encrypts files The PUT command in Snowflake compresses files using Gzip by default and encrypts files during the transfer to a stage, ensuring secure and efficient data loading.

upvoted 2 times

**JasMozai** 2 years, 6 months ago

PUT does not support uploading files to external stages. All files stored on internal stages for data loading and unloading operations are automatically encrypted using AES-256 strong encryption on the server side. By default, Snowflake provides additional client-side encryption with a 128-bit key (with the option to configure a 256-bit key). AUTO_COMPRESS is TRUE by default and by default it uses GZIP.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

c and d

upvoted 1 times

**BobFar** 3 years ago

CD are the correct answer!

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

CD are the correct options

upvoted 1 times

**ramisou** 3 years, 7 months ago

CD are the correct options

upvoted 3 times

**ajay_1233456** 3 years, 9 months ago

CD are the correct options

upvoted 3 times

**N89** 3 years, 10 months ago

CD are the correct options

upvoted 3 times

**SV1122** 3 years, 10 months ago

Uploaded files are automatically encrypted with 128-bit or 256-bit keys. AUTO_COMPRESS is TRUE by default and by default it uses GZIP.

https://docs.snowflake.com/en/sql-reference/sql/put.html

upvoted 4 times

**Tapasgup007** 3 years, 11 months ago

CD are correct

upvoted 1 times

**bharathmv** 4 years, 1 month ago

CD is correct

upvoted 2 times

**bharathmv** 4 years, 1 month ago

CD is right answer

upvoted 1 times

**Percy2112** 4 years, 2 months ago

CD are default in snowflake. For compression and encryption , snowflake take care. AB are the correct option.

upvoted 1 times

Question #28 *Topic 1*

Which type of table corresponds to a single Snowflake session?

A. Temporary

B. Transient

C. Provisional

D. Permanent

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**KaranSnow** Most Recent 10 months, 1 week ago

Temporary table only meant for single session

upvoted 2 times

**jeny12138** 1 year, 7 months ago

temporary is the signle

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

A.Temporary

upvoted 4 times

**ASinanan** 2 years, 7 months ago

A is correct. https://docs.snowflake.com/en/user-guide/tables-temp-transient

upvoted 4 times

**_yyukta** 2 years, 8 months ago

A is correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**RajuNaik** 3 years, 6 months ago

A is correct

upvoted 1 times

**Ayush24** 3 years, 6 months ago

Temporary tables are created in a session and they cannot be accessed through different session.

upvoted 2 times

**alanhuangjob** 3 years, 7 months ago

Answer is Temporary table.

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

A is correct

upvoted 1 times

**sdnv89** 3 years, 11 months ago

Temporary tables only exist within the session in which they were created and persist only for the remainder of the session.

https://docs.snowflake.com/en/user-guide/tables-temp-transient.html#temporary-tables

upvoted 4 times

**Percy2112** 4 years, 2 months ago

Temporary Tables.

upvoted 2 times

**Sundarr** 4 years, 9 months ago

Temporary table.

upvoted 4 times

Question #29 *Topic 1*

Which interfaces can be used to create and/or manage Virtual Warehouses?

A. The Snowflake Web Interface (UI)

B. SQL commands

C. Data integration tools

D. All of the above

**Suggested Answer:** *D*

*Community vote distribution*

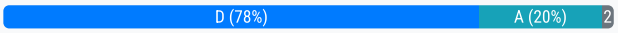


**Graghu** Highly Voted 4 years, 6 months ago

https://docs.snowflake.com/en/user-guide/intro-key-concepts.html A warehouse can be created through the web interface or using SQL: Classic Web Interface Click on Warehouses Warehouses tab » Create SQL Execute a CREATE WAREHOUSE command.

upvoted 14 times

**JoyCao** Most Recent 9 months, 2 weeks ago

Google AI Answer: many third-party data integration tools that connect to Snowflake (such as Talend, Informatica, SnapLogic, or ActiveBatch) are capable of creating and managing virtual warehouses, often by executing SQL commands like CREATE WAREHOUSE on the user's behalf.

upvoted 1 times

**KaranSnow** 10 months, 1 week ago

It's D. Data integration tools: Various third-party data integration and ETL tools (like Talend, Matillion, Fivetran, Pentaho) can interact with Snowflake and automate warehouse management tasks, often by executing SQL commands or using dedicated connectors under the hood, as part of data processing workflows or data warehouse automation.

upvoted 1 times

**shreyashichak** 1 year, 1 month ago

A and B. Integration tools are external to snowflake. How can they be used to create WH?

upvoted 1 times

**visionary4ever** 1 year, 2 months ago

Data integration tools are Informatica, Talend, Matillion, Fivetran

upvoted 1 times

**mishrask06** 1 year, 2 months ago

D is wrong, not sure why this is voted highest. You cannot create warehouse through any integration tool.

upvoted 3 times

**jeny12138** 1 year, 7 months ago

D is correct

upvoted 1 times

**4762448** 1 year, 7 months ago

" In addition, the scope of data integration has expanded to include a wider range of operations, including: - Data preparation. - Data migration, movement, and management. - Data warehouse automation.* " https://docs.snowflake.com/en/user-guide/ecosystem-etl Taking into account the last *Data warehouse automation*, I think the correct answer is D

upvoted 1 times

**XXXXmas** 1 year, 10 months ago

All warehouse tasks can be performed from the Snowflake web interface or using the DDL commands for warehouses.
https://docs.snowflake.com/en/user-guide/warehouses-tasks

upvoted 1 times

**krishna3195** 2 years ago

A. The Snowflake Web Interface (UI): True. Snowflake provides a user-friendly web interface (UI) where users can create, resize, suspend, resume, and manage virtual warehouses. B. SQL commands: True. Virtual warehouses can be created and managed using SQL commands like CREATE WAREHOUSE, ALTER WAREHOUSE, DROP WAREHOUSE, etc. C. Data integration tools: True. Many data integration tools that connect to Snowflake (such as ETL tools like Talend, Matillion, or Fivetran) can also manage virtual warehouses, either through direct integration or by executing SQL commands under the hood. Since virtual warehouses can be managed via all of the listed interfaces, the correct answer is D. All of the above.

upvoted 4 times

**d22770a** 2 years, 1 month ago

Everyone who believes D is right answer, can you please "NAME" the data integration tool with which we can create a Virtual WH. The option says DATA INTEGRATION "TOOLS" , so I am assuming if this is correct then there should be more than 1. Please help

upvoted 4 times

**abhijitrajan123** 2 years, 1 month ago

A and B

upvoted 2 times

**aemilka** 2 years, 2 months ago

In addition, the scope of data integration has expanded to include a wider range of operations, including: - Data preparation. - Data migration, movement, and management. - Data warehouse automation. https://docs.snowflake.com/en/user-guide/ecosystem-etl Data integration includes data warehouse automation, so D is correct.no, Its talking about data warehouse while question about virtual warehouse

upvoted 1 times

**aemilka** 2 years, 2 months ago

**Selected Answer: D**

In addition, the scope of data integration has expanded to include a wider range of operations, including: - Data preparation. - Data migration, movement, and management. - Data warehouse automation.

https://docs.snowflake.com/en/user-guide/ecosystem-etl Data integration includes data warehouse automation, so D is correct.

upvoted 1 times

**sachindbasoni** 2 years, 1 month ago

no, Its talking about data warehouse while question about virtual warehouse

upvoted 3 times

**ABRocks** 2 years, 2 months ago

A and B

upvoted 1 times

**VLMBL_94** 2 years, 3 months ago

Ans : A,B

upvoted 1 times

**vivekmani2021** 2 years, 3 months ago

A and B , Data integration tolls cannot create VW

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

A AND B

upvoted 1 times

Question #30 *Topic 1*

When a Pipe is recreated using the CREATE OR REPLACE PIPE command:

A. The Pipe load history is reset to empty

B. The REFRESH parameter is set to TRUE

C. Previously loaded files will be ignored

D. All of the above

**Suggested Answer:** *A*

*Community vote distribution*

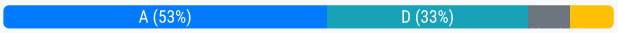


**jjordan** Highly Voted 4 years, 10 months ago

I believe the correct answer is A. REFRESH is a parameter for ALTER PIPE https://docs.snowflake.com/en/sql-reference/sql/alter-pipe.html It is NOT a parameter for CREATE [OR REPLACE] Pipe. https://docs.snowflake.com/en/sql-reference/sql/create-pipe.html Further recreating a pipe resets history: https://docs.snowflake.com/en/user-guide/data-load-snowpipe-manage.html#label-snowpipe-management-recreate-pipes "Load History The load history for Snowpipe operations is stored in the metadata of the pipe object. When a pipe is recreated, the load history is dropped. In general, this condition only affects users if they subsequently execute an ALTER PIPE … REFRESH statement on the pipe. Doing so could load duplicate data from staged files in the storage location for the pipe if the data was already loaded successfully and the files were not deleted subsequently."

upvoted 28 times

**KaranSnow** 10 months, 1 week ago

Answer is D.

upvoted 3 times

**ujkim** 1 year ago

Snowflake Docs: Managing Pipes — Recreate Pipes The load history for Snowpipe operations is stored in the metadata of the pipe object. When a pipe is recreated, the load history is dropped. Snowflake Docs: ALTER PIPE

upvoted 2 times

**ShawnZh** 1 year, 1 month ago

B, C have not Snowflake document to support that, ChatGPT shows D, but after deeper asks, CHATGPT shows that it has no document support options B and C.

upvoted 1 times

**niho1993** 1 year, 1 month ago

D. All of the above the correct answer

upvoted 2 times

**iyiola_daniel** 1 year, 1 month ago

Answer from Gemini When a Pipe is recreated using the CREATE OR REPLACE PIPE command in Snowflake, the following occurs: The Pipe load history is reset to empty. This means the pipe will no longer track files that were previously loaded by the old pipe definition. The REFRESH parameter is automatically set to TRUE. This ensures that the newly created pipe will immediately begin looking for new files in the specified stage location. Previously loaded files will be ignored. Since the load history is reset and REFRESH is enabled, the pipe will not reprocess files that were already successfully loaded by the previous pipe definition. Therefore, the correct answer is D. All of the above.

upvoted 3 times

**Sk_3** 1 year, 1 month ago

D is correct answer

upvoted 1 times

**suhaxini1** 1 year, 2 months ago

it seems that none is correctD is correct, for official doc: https://docs.snowflake.com/en/user-guide/data-load-snowpipe-manage#label-snowpipe-management-recreate-pipes Load history The load history for Snowpipe operations is stored in the metadata of the pipe object. When a pipe is recreated, the load history is dropped. In general, this condition only affects users if they subsequently execute an ALTER PIPE … REFRESH statement on the pipe. Doing so could load duplicate data from staged files in the storage location for the pipe if the data was already loaded successfully and the files were not deleted subsequently. This means that the " Previously loaded files will be ignored" if you after "CREATE OR REPLACE PIPE" don't explic execute "ALTER PIPE … REFRESH".

upvoted 1 times

**suhaxini1** 1 year, 2 months ago

**Selected Answer: A**

it seems that none is correct

upvoted 1 times

**BigDataBB** Most Recent 4 months, 1 week ago

D is correct, for official doc: https://docs.snowflake.com/en/user-guide/data-load-snowpipe-manage#label-snowpipe-management-recreate-pipes Load history The load history for Snowpipe operations is stored in the metadata of the pipe object. When a pipe is recreated, the load history is dropped. In general, this condition only affects users if they subsequently execute an ALTER PIPE … REFRESH statement on the pipe. Doing so could load duplicate data from staged files in the storage location for the pipe if the data was already loaded successfully and the files were not deleted subsequently. This means that the " Previously loaded files will be ignored" if you after "CREATE OR REPLACE PIPE" don't explic execute "ALTER PIPE … REFRESH".

upvoted 1 times

**lm_bp** 1 year, 2 months ago

When a Pipe is recreated using the CREATE OR REPLACE PIPE command in Snowflake, the following happens: A. The Pipe load history is reset to empty True: When you recreate a pipe, Snowflake resets the load history associated with that pipe. This means that the system no longer has any record of files that have been previously loaded through that pipe. B. The REFRESH parameter is set to TRUE True: When you recreate the pipe, the REFRESH parameter is set to TRUE by default. This ensures that Snowflake will reprocess the files that are staged for loading, allowing the pipe to refresh its state and start processing new data again. C. Previously loaded files will be ignored True: After the CREATE OR REPLACE PIPE operation, previously loaded files (files that were ingested before the pipe was replaced) will be ignored. This is because the pipe's state is reset, and Snowflake does not automatically track the files that have been processed before the pipe was replaced.

upvoted 2 times

**slushmush** 1 year, 5 months ago

The answers are A and C, but the question only allows one option to be selected. Previously loaded files will be ignored unless an ALTER PIPE... REFRESH is executed: https://docs.snowflake.com/en/user-guide/data-load-snowpipe-manage#label-snowpipe-management-recreate-pipes Refresh is not an option when creating a pipe: https://docs.snowflake.com/en/sql-reference/sql/create-pipe

upvoted 1 times

**Vijaya15** 1 year, 6 months ago

D Snowflake clears the load history for the pipe. Setting REFRESH = TRUE tells Snowflake to re-evaluate the staged files. Snowpipe uses the new pipe definition and skips previously processed files,

upvoted 1 times

**jeny12138** 1 year, 7 months ago

A is correct

upvoted 1 times

**4762448** 1 year, 7 months ago

Correct answer is A,C

upvoted 2 times

**PrasantSadangi** 1 year, 7 months ago

Answer is A The pipe load history is set to empty. The load history for Snowpipe operations is stored in the metadata of the pipe object. When a pipe is recreated, the load history is dropped.In general, this condition only affects users if they subsequently execute an ALTER PIPE … REFRESH statement on the pipe. Doing so could load duplicate data from staged files in the storage location for the pipe if the data was already loaded successfully and the files were not deleted subsequently.

upvoted 1 times

**akomal** 1 year, 7 months ago

when recreating a pipe with CREATE OR REPLACE PIPE, all these behaviors apply, making D. All of the above the correct answer

upvoted 1 times

**Cartoffee** 1 year, 8 months ago

D is correct here

upvoted 1 times

**kentucky13caa** 1 year, 11 months ago

The correct answer is D. All of the above. When a Pipe is recreated using the CREATE OR REPLACE PIPE command: The Pipe load history is reset to empty (A is true) The REFRESH parameter is automatically set to TRUE (B is true) Previously loaded files will be ignored (C is true) Therefore, since all three statements (A, B, and C) are true, D "All of the above" is the correct answer. This is important to understand because recreating a pipe effectively gives you a fresh start - clearing history, refreshing the pipe, and ignoring previously processed files.

upvoted 2 times

Question #31 *Topic 1*

What is the minimum Snowflake edition that customers planning on storing protected information in Snowflake should consider for regulatory compliance?

A. Standard

B. Premier

C. Enterprise

D. Business Critical Edition

**Suggested Answer:** *D*

*Community vote distribution*

D (55%) C (45%)

**jjordan** Highly Voted 4 years, 10 months ago

The question should be more specific about the type(s) of "protected information" they are talking about. For some "protected data" SOC2 compliance is adequate.

upvoted 8 times

**kentucky13caa** Highly Voted 1 year, 11 months ago

The correct answer is C. Enterprise The Enterprise edition is the minimum Snowflake edition that customers should consider when planning to store protected information requiring regulatory compliance. This is because the Enterprise edition includes essential security and compliance features such as: 1. Enhanced security features 2. Advanced data protection capabilities 3. Compliance-related functionalities 4. More robust governance controls The Standard (A) and Premier (B) editions don't provide sufficient compliance and security features for protected information, while Business Critical Edition (D) offers even more advanced features but isn't the minimum required edition for regulatory compliance - Enterprise is.

upvoted 5 times

**Wzghaer** Most Recent 10 months, 1 week ago

D is the right answer : Business Critical Edition, formerly known as Enterprise for Sensitive Data (ESD), offers even higher levels of data protection to support the needs of organizations with extremely sensitive data, particularly PHI data that must comply with HIPAA and HITRUST CSF regulations. HITRUST CSF Certification This certification enhances Snowflake's security posture in regulatory compliance and risk management, and applies to Snowflake editions that are Business Critical (or higher). For more information, see Snowflake Security and Trust Center.

upvoted 1 times

**nizmo63** 1 year, 3 months ago

C is the right one as its talking about minimum

upvoted 1 times

**98198d9** 1 year, 5 months ago

C is the correct one when we asking for minimum

upvoted 1 times

**jeny12138** 1 year, 7 months ago

BCE is for regulary

upvoted 1 times

**BrunoRYamada** 1 year, 7 months ago

is easy to think of D, but MINIMUM is keyword, so C already fits

upvoted 1 times

**PrasantSadangi** 1 year, 7 months ago

Minimum is the key

upvoted 1 times

**Cartoffee** 1 year, 8 months ago

D would be better as it has more features but questions asks for minimum - C is correct here

upvoted 1 times

**XXXXmas** 1 year, 10 months ago

https://docs.snowflake.com/en/user-guide/intro-editions#business-critical-edition

upvoted 1 times

**jiriz** 2 years, 1 month ago

I think, it's Enterprise, so C. Business Critical Edition (D) is definitely better choice and recommended, but the question is about "minimum".

upvoted 4 times

**Mallikharjuna452** 2 years, 3 months ago

D is Corect Answer

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D IS CORRECT

upvoted 3 times

**shubtred** 3 years, 4 months ago

C. Enterprise Edition. The Enterprise Edition of Snowflake offers advanced security features and capabilities that are designed to meet the stringent requirements of regulatory compliance. It includes features such as role-based access control (RBAC), encryption at rest and in transit, audit logging, fine-grained access controls, and data masking. These features help ensure the security and compliance of sensitive and protected data stored in Snowflake. While other editions of Snowflake, such as Standard and Premier, also provide security features, the Enterprise Edition is specifically tailored for organizations with higher compliance and security requirements. It offers additional features and capabilities that address the needs of regulated industries and data protection standards. Therefore, customers planning on storing protected information in Snowflake for regulatory compliance should consider the Enterprise Edition as a suitable option.

upvoted 4 times

**Mayuri917233** 3 years, 5 months ago

D is correct

upvoted 1 times

**dotaho** 3 years, 9 months ago

D is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

D is correct

upvoted 1 times

Question #32 *Topic 1*

Select the three types of tables that exist within Snowflake. (Choose three.)

A. Temporary

B. Transient

C. Provisional

D. Permanent

**Suggested Answer:** *A, B, D*

*Community vote distribution*

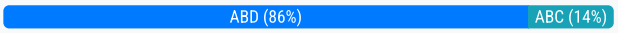


**jeny12138** 1 year, 7 months ago

ABC is correctdon't exists an Provisional Table in Snowflake

upvoted 1 times

**jeny12138** 1 year, 7 months ago

**Selected Answer: ABC**

ABC is correct

upvoted 1 times

**BigDataBB** Most Recent 4 months, 1 week ago

don't exists an Provisional Table in Snowflake

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

ABD ARE CORRECT

upvoted 1 times

**_yyukta** 2 years, 8 months ago

ABD ARE CORRECT

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

ABD are correct

upvoted 1 times

**Ayush24** 3 years, 6 months ago

Temporary tables, Transient tables, Permanent tables exists in Snowflake.

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

ABD is the correct answer

upvoted 3 times

**Diegovaz** 4 years, 4 months ago

There is external too

upvoted 1 times

**rikku33** 4 years, 5 months ago

A. Temporary B. Transient D. Permanent

upvoted 1 times

**moon2351** 4 years, 10 months ago

Answer is correct.

upvoted 1 times

Question #33 *Topic 1*

True or False: Snowpipe via REST API can only reference External Stages as source.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*




**Tapsinha** Highly Voted 4 years, 9 months ago

Statement is False. Answer B is correct. Snowpipe via rest API can refer both external as well as internal stages.
https://docs.snowflake.com/en/user-guide/data-load-snowpipe-rest-overview.html

upvoted 16 times

**VITO7** Most Recent 1 year, 6 months ago

Statement is False. Answer B is correct.

upvoted 2 times

**jeny12138** 1 year, 7 months ago

false for only external

upvoted 1 times

**Yugendharsai** 1 year, 10 months ago

Supports both Internal and External Internal is Snowflake and External is S3 Blob etc

upvoted 1 times

**kentucky13caa** 1 year, 11 months ago

The correct answer is A. True Snowpipe via REST API can only reference External Stages as source. This is a key limitation of using Snowpipe with the REST API. Internal stages cannot be used as a source when using Snowpipe through the REST API interface - you must use external stages. This is an important distinction from other Snowpipe usage patterns.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B.False

upvoted 1 times

**deeshubh** 2 years, 6 months ago

This statement is true because Snowflake, as a cloud-based data warehouse, requires a specific driver to establish a connection. While third-party tools that support standard JDBC (Java Database Connectivity) or ODBC (Open Database Connectivity) protocols theoretically could connect to Snowflake, without a Snowflake-specific driver, they would not be able to establish a connection successfully. Snowflake's architecture and features require a driver that is specifically designed to work with its unique system. Therefore, a third-party tool lacking a Snowflake-specific driver would indeed be unable to connect to Snowflake.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B is correct

upvoted 1 times

**lischris** 3 years, 4 months ago

False: https://docs.snowflake.com/en/user-guide/data-load-snowpipe-rest-gs#step-4-stage-data-files Snowpipe supports loading from the following stage types: Named internal (Snowflake) or external (Amazon S3, Google Cloud Storage, or Microsoft Azure) stages Table stages

upvoted 1 times

**snehav112** 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/data-load-snowpipe-rest-load B) Stage your data files: Internal stage: Use the PUT command to stage your files. External stage: Use the client tools provided by the cloud provider to copy your files to the stage location (Amazon S3, Google Cloud Storage, or Microsoft Azure).

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**sidyndiaye** 3 years, 7 months ago

Answer will be false

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

B is correct

upvoted 1 times

**learnazureportal** 4 years, 5 months ago

The correct answer is B.

upvoted 2 times

**San2** 4 years, 6 months ago

REST API supports both internal and external but SNOWPIPE AUTO INGEST only support external. so the answer is B

upvoted 1 times

**StrStudy** 4 years, 9 months ago

answer is false

upvoted 3 times

**mayur1234** 4 years, 9 months ago

Answer is B. Snowpipe via rest API can refer both external as well as internal stages. https://docs.snowflake.com/en/user-guide/data-load-snowpipe-rest-overview.html

upvoted 1 times

Question #34 *Topic 1*

True or False: A third-party tool that supports standard JDBC or ODBC but has no Snowflake-specific driver will be unable to connect to Snowflake.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*




**redsky0** Highly Voted 4 years, 10 months ago

Answer is B False: E.g. JDBC: https://docs.snowflake.com/en/user-guide/jdbc.html Snowflake provides a JDBC type 4 driver that supports core JDBC functionality. The JDBC driver must be installed in a 64-bit environment and requires Java 1.8 (or higher). The driver can be used with most client tools/applications that support JDBC for connecting to a database server.

upvoted 29 times

**shubtred** Highly Voted 3 years, 4 months ago

False. A third-party tool that supports standard JDBC or ODBC can connect to Snowflake even without a Snowflake-specific driver. Snowflake provides JDBC and ODBC connectors that adhere to the standard JDBC and ODBC interfaces, allowing third-party tools to establish a connection to Snowflake using these standard protocols. The Snowflake JDBC and ODBC connectors act as bridges between the third-party tools and the Snowflake service, enabling communication and data exchange. As long as the third-party tool supports JDBC or ODBC, it can utilize the Snowflake connectors to connect to Snowflake and interact with the data warehouse. This approach ensures compatibility and interoperability with a wide range of tools and applications that support JDBC or ODBC, making it easier to integrate Snowflake into existing ecosystems or use preferred third-party tools for data analysis, reporting, or other purposes.

upvoted 14 times

**Denisa70** Most Recent 9 months, 4 weeks ago

FALSE no Snowflake specific driver needed for any of the connectors... thats why we need connectors in the first place

upvoted 1 times

**niho1993** 1 year, 1 month ago

a table can be queried using any Virtual Warehouse, not just the one that loaded the data.

upvoted 1 times

**nizmo63** 1 year, 3 months ago

yes it is possible to not select a file format

upvoted 1 times

**nizmo63** 1 year, 3 months ago

SF has driver support for most connectors

upvoted 1 times

**user_1011** 1 year, 5 months ago

Answer is False: It is mentioned in the official snowflake documentation to use net.snowflake.client.jdbc.SnowflakeDriver as the driver class in your JDBC application. So if there are no snowflake specific driver, even through JDBC, it is not possible to establish the connection in snowflake https://docs.snowflake.com/en/developer-guide/jdbc/jdbc-configure

upvoted 1 times

**Sudhansu21** 1 year, 8 months ago

A third-party tool that supports standard JDBC or ODBC can connect to Snowflake even if it does not have a Snowflake-specific driver. Snowflake provides its own JDBC and ODBC drivers that can be used by any tool that supports these standard interfaces. Therefore, as long as the tool can use a JDBC or ODBC driver, it can connect to Snowflake using Snowflake's provided drivers.

upvoted 1 times

**rocky_vits1** 1 year, 10 months ago

JDBC Driver is provided by snowflake

upvoted 2 times

**legohax** 1 year, 10 months ago

False, Snowflake has generic ODBC/JDBC driver support

upvoted 2 times

**Yugendharsai** 1 year, 10 months ago

Snowflake supports JDBC OR ODBC Does not care about driver it having its supported driver version to connect

upvoted 1 times

**Djama** 1 year, 11 months ago

Answer is B : False

upvoted 1 times

**kentucky13caa** 1 year, 11 months ago

The correct answer is B. False A third-party tool that supports standard JDBC or ODBC can connect to Snowflake even without a Snowflake-specific driver. This is because Snowflake supports standard JDBC and ODBC protocols, allowing tools that use these standard database connectivity interfaces to establish connections to Snowflake. The statement is false because: 1. Standard JDBC/ODBC support is sufficient for basic connectivity 2. Snowflake is designed to work with standard database connectivity protocols 3. While Snowflake-specific drivers might offer optimized performance or additional features, they are not strictly required for basic connectivity

upvoted 1 times

**Shipra123** 1 year, 11 months ago

False. A third-party tool that supports standard JDBC or ODBC can connect to Snowflake even without a Snowflake-specific driver.

upvoted 1 times

**HICH9173** 2 years ago

Answer is B

upvoted 1 times

**pjfunner** 2 years ago

It's B. As long as ODBC and JDBC are supported you can connect to Snowflake.

upvoted 1 times

**Daniel1412** 2 years ago

as long as JDBC or ODBC is supported it can connect

upvoted 1 times

Question #35 *Topic 1*

True or False: It is possible to load data into Snowflake without creating a named File Format object.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*




**Mallikharjuna452** Most Recent 2 years, 3 months ago

Yes It is possible,A.True

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. True

upvoted 1 times

**Colourseun** 3 years ago

A: True is correct Snowflake loads data as CSV by default, but if the file is not CSV, you will have to specify the File Format.

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**SteveCrown** 3 years, 6 months ago

True. Snowflake supports loading data without creating a named File Format object. Snowflake provides built-in file formats that can be used to load data without the need to create a file format object explicitly. For example, to load a CSV file, you can use the built-in CSV file format by simply specifying the FILE_FORMAT option as CSV in the COPY INTO command. This allows you to quickly load data without the overhead of creating a named file format object. Send a message...

upvoted 2 times

**abbcia** 3 years, 7 months ago

a is true, you can specify a format but you don't have to

upvoted 2 times

**Raju039** 3 years, 8 months ago

https://docs.snowflake.com/en/user-guide/data-load-external-tutorial-create-file-format

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

A is correct

upvoted 3 times

**SD3** 4 years, 5 months ago

As the link suggests: 'In regular use, this step is optional, but is recommended'. Therefore it can be loaded without. I think thats enough proof?

upvoted 4 times

**Graghu** 4 years, 6 months ago

table will be loaded with CSV format as default format if there is no file format defined

upvoted 3 times

**yayaro** 4 years, 7 months ago

Needs a proof that shows A is correct. Until then I will go with B.

upvoted 2 times

**arpit_dataguy** 4 years, 9 months ago

A is the correct answer. With a table created we can upload data from a csv file via web UI directly in the table without defining the FILE FORMAT.

upvoted 3 times

Question #36 *Topic 1*

True or False: A table in Snowflake can only be queried using the Virtual Warehouse that was used to load the data.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Tapwal** Most Recent 1 year, 10 months ago

In Snowflake, tables can be queried using any Virtual Warehouse within the same account, not just the one used to load the data. This flexibility is due to Snowflake's architecture, which separates storage and compute resources.

upvoted 1 times

**56d43d6** 1 year, 10 months ago

while this is false, I don't think the question is asking about whether you can get to the result cache. I think the point is, access to databases, schemas and tables is managed via grants.

upvoted 1 times

**Yosuke** 1 year, 11 months ago

B is correct. We can get data from the result cache in the cloud service layer, without a virtual warehouse.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

b.False

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B.False

upvoted 1 times

**marklv** 3 years, 1 month ago

a multi cluster, any warehouse can get the data

upvoted 1 times

**douglaszickuhr** 3 years, 5 months ago

B for sure

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**tarekff** 3 years, 6 months ago

B is correct. you can query cached data without having virtual warehouse running

upvoted 2 times

**alanhuangjob** 3 years, 7 months ago

B is correct

upvoted 1 times

**EmiB** 3 years, 7 months ago

B - correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

B is correct

upvoted 2 times

**Percy2112** 4 years, 2 months ago

No Doubt . B is correct.

upvoted 1 times

**Sundarr** 4 years, 9 months ago

Answer is B

upvoted 1 times

**Zafar202** 4 years, 9 months ago

Ofcourse B :)

upvoted 1 times

**moon2351** 4 years, 10 months ago

Answer is B

upvoted 1 times

Question #37 *Topic 1*

Which of the following statements are true of Snowflake data loading? (Choose three.)

A. VARIANT ɔ€nullɔ€ values are not the same as SQL NULL values

B. It is recommended to do frequent, single row DMLs

C. It is recommended to validate the data before loading into the Snowflake target table

D. It is recommended to use staging tables to manage MERGE statements

**Suggested Answer:** *A, C, D*

*Community vote distribution*

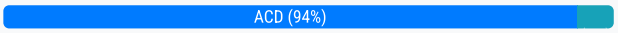


**BrunoRYamada** Most Recent 1 year, 7 months ago

B talks about refrequent single row DMLs, which goes against columnar principle

upvoted 1 times

**Sudhansu21** 1 year, 8 months ago

A. VARIANT "null" values are not the same as SQL NULL values: In Snowflake, a VARIANT type can store semi-structured data like JSON. A JSON null is distinct from an SQL NULL. JSON null is treated as a valid value in Snowflake, while SQL NULL represents the absence of a value. B. It is recommended to do frequent, single row DMLs: This is false. Snowflake is optimized for bulk operations, and frequent single-row DMLs (e.g., INSERT, UPDATE, DELETE) are not recommended as they can lead to inefficiencies and higher costs. C. It is recommended to validate the data before loading into the Snowflake target table: This is true. Validating data before loading ensures data quality and avoids errors or inconsistencies in the target table. D. It is recommended to use staging tables to manage MERGE statements: This is true. Staging tables are often used to prepare and validate data before applying MERGE operations, which helps ensure data integrity and efficiency. Correct Choices: A, C, D

upvoted 4 times

**Mallikharjuna452** 2 years, 3 months ago

ACD are correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

ACD are correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

ACD are correct

upvoted 3 times

**dmitriypo** 3 years, 6 months ago

Variant NULL values: https://docs.snowflake.com/en/user-guide/semistructured-considerations#label-variant-null

upvoted 2 times

**alanhuangjob** 3 years, 7 months ago

ACD is correct answer.

upvoted 1 times

**KarBiswa** 3 years, 9 months ago

The result must be deterministic for merge so source can be any reference data source only constraint it should create/modify unique rows

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

ACD are correct

upvoted 1 times

**Antz83** 4 years, 3 months ago

Frequent Single rows on any MPP like Snowflake has always been a no no! Function of an OLTP DB

upvoted 3 times

**MinuteMaid** 4 years, 7 months ago

Frequent single row DML? Hell nah. ACD is the way to go.

upvoted 3 times

**Rajokkiyam** 4 years, 9 months ago

ACD, Correct

upvoted 2 times

Question #38 *Topic 1*

Which statements are true of micro-partitions? (Choose two.)

A. They are approximately 16MB in size

B. They are stored compressed only if COMPRESS=TRUE on Table

C. They are immutable

D. They are only encrypted in the Enterprise edition and above

**Suggested Answer:** *A, C*

*Community vote distribution*

AC (100%)

**EmiB** Highly Voted 3 years, 8 months ago

Selected Answer: AC

- micro-partitions are small in size (50 to 500 MB, before compression) - 16 MB must comes from old version/question - COMPRESSION = NONE is the right

upvoted 7 times

**EmiB** Highly Voted 3 years, 8 months ago

- micro-partitions are small in size (50 to 500 MB, before compression) - 16 MB must comes from old version/question - COMPRESSION = NONE is the rightYour answer was a little unclear to me. Uncompressed the data of a micro partition would be around 50-500 MB, while compressed it is around 16 MB. Second answer is that these micro partitions are immutable

upvoted 7 times

**TheHuman_** 2 years, 6 months ago

Your answer was a little unclear to me. Uncompressed the data of a micro partition would be around 50-500 MB, while compressed it is around 16 MB. Second answer is that these micro partitions are immutable

upvoted 2 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A and C are Correct.after loading into the table the data would be compressed.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A and C are correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

AC are correct

upvoted 1 times

**alanhuangjob** 3 years, 7 months ago

I believe A and C are correct.

upvoted 1 times

**SV1122** 3 years, 10 months ago

The size of 50-500MB is for uncompressed data and micropartition itself holds around 16MB(after compression).
https://stackoverflow.com/questions/67945427/clarification-on-the-snowflake-micro-partition-size

upvoted 3 times

**Tapasgup007** 3 years, 11 months ago

AC are correct

upvoted 1 times

**Rajokkiyam** 4 years, 9 months ago

AC correct

upvoted 2 times

**Sundarr** 4 years, 9 months ago

A and C. They are approximately 16MB in size and also immutable.

upvoted 1 times

**arpit_dataguy** 4 years, 9 months ago

16 MB compressed

upvoted 2 times

**moon2351** 4 years, 10 months ago

Answer is correct

upvoted 3 times

Question #39 *Topic 1*

True or False: Query ID's are unique across all Snowflake deployments and can be used in communication with Snowflake Support to help troubleshoot issues.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*




**Anirudh2020** Highly Voted 4 years, 8 months ago

A is correct

upvoted 17 times

**Lena_Frazier** Highly Voted 2 years, 9 months ago

https://community.snowflake.com/s/question/0D5Do000011qRx6KAE/is-query-ids-are-unique-across-all-snowflake-deployments-and-can-be-used-in-communication-with-snowflake-support-to-help-troubleshoot-issues

upvoted 7 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Yes True

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct answer:A

upvoted 2 times

**Stevec80** 2 years, 5 months ago

Snowflake Query_ID is generated by UUID and is UNIQUE across all the snowflake accounts

upvoted 3 times

**_yyukta** 2 years, 8 months ago

B. False

upvoted 1 times

**VarnikaSri** 2 years, 9 months ago

True Snowflake Query_ID is generated by UUID and is UNIQUE across all regions

https://community.snowflake.com/s/question/0D5Do000011qRx6KAE/is-query-ids-are-unique-across-all-snowflake-deployments-and-can-be-used-in-communication-with-snowflake-support-to-help-troubleshoot-issues

upvoted 1 times

**avashish** 2 years, 9 months ago

Query ID should be unique across all sessions. A is correct answer

upvoted 2 times

**suhovet** 2 years, 10 months ago

A is correct

upvoted 1 times

**6c79d6f** 2 years, 11 months ago

Answer would be A in accordance with the following: https://community.snowflake.com/s/article/Check-for-a-query-executed-in-the-session

upvoted 1 times

**zephyrvz** 3 years ago

A is correct

upvoted 2 times

**Gwo** 3 years, 1 month ago

A is true. QueryID is based on UUID and they are unique among all the snowflake accounts. So when you are submitting a case and you just provide queryID and nothing else even in those cases, Snowflake support will be able to locate your query. Though Support may obviously ask for more information in order to help you better.

upvoted 2 times

**Marge23** 3 years, 1 month ago

TRUE, you can send your QID to SF support to help solve issues.

upvoted 2 times

**Marge23** 3 years, 2 months ago

TRUE you need to send your QID to the SF support to evaluate errors too.

upvoted 2 times

**lofibeats** 3 years, 2 months ago

it's true

upvoted 1 times

**PoornaChandraIn** 3 years, 3 months ago

Answer is A. QueryID is based on UUID and they are unique among all the snowflake accounts.

upvoted 4 times

**Snow_P** 3 years, 4 months ago

True: Upon execution each query in Snowflake receives its' own unique query ID which can be used to access the query results at a later stage, obtain help on the query from Snowflake support or try to analyze or tune the query execution. https://community.snowflake.com/s/article/Check-for-a-query-executed-in-the-session

upvoted 2 times

Question #40 *Topic 1*

A deterministic query is run at 8am, takes 5 minutes, and the results are cached. Which of the following statements are true? (Choose two.)

A. The exact query will ALWAYS return the precomputed result set for the RESULT_CACHE_ACTIVE = time period

B. The same exact query will return the precomputed results if the underlying data hasn't changed and the results were last accessed within previous 24 hour period

C. The same exact query will return the precomputed results even if the underlying data has changed as long as the results were last accessed within the previous 24 hour period

D. The €24ı hourɜ€ timer on the precomputed results gets renewed every time the exact query is executed

**Suggested Answer:** *B, D*

*Community vote distribution*

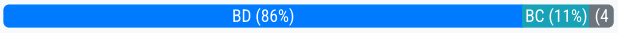


**ggww** Highly Voted 4 years, 10 months ago
correct answer is B, D
upvoted 9 times

**4762448** Most Recent 1 year, 7 months ago
b and d are the correct options
upvoted 1 times

**BrunoRYamada** 1 year, 7 months ago
B & D, cache is only for 24hr, if data doesn't change, and renews TTL if queried again
upvoted 1 times

**Sudhansu21** 1 year, 8 months ago
B -Snowflake's result cache will return the precomputed results if the data has not changed and the cached results are still valid (within 24 hours of the last access). D- Every time the query is executed, the 24-hour timer for the cached results is reset, extending the validity of the cached results.
upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago
B and D
upvoted 1 times

**_yyukta** 2 years, 8 months ago
B and D are correct
upvoted 1 times

**444823b** 2 years, 10 months ago
esas son las correctas b y d
upvoted 1 times

**rich922** 3 years, 4 months ago
BD I think, but bad wording on B: "previous" 24 hours would mean before the original query was run, so taken literally this is incorrect but I think they mean the next 24 hours.
upvoted 1 times

**Mayuri917233** 3 years, 5 months ago
correct answer is B, D
upvoted 1 times

**SV1122** 3 years, 10 months ago
BD Explanation for D is here https://docs.snowflake.com/en/user-guide/querying-persisted-results.html
upvoted 3 times

**Tapasgup007** 3 years, 11 months ago
BD are correct
upvoted 1 times

**Graghu** 4 years, 6 months ago

BD is the correct answer

upvoted 2 times

**Sly007** 4 years, 6 months ago

BD for sure

upvoted 3 times

**MinuteMaid** 4 years, 7 months ago

Obviously it is B,D by MinuteMaid

upvoted 1 times

**Fosterchild123** 4 years, 8 months ago

The results cache will only be used if it's the same query AND the same role executing the query

upvoted 1 times

**Anirudh2020** 4 years, 8 months ago

B and D are correct

upvoted 1 times

**Rajokkiyam** 4 years, 9 months ago

answer BD

upvoted 1 times

Question #41 *Topic 1*

Increasing the maximum number of clusters in a Multi-Cluster Warehouse is an example of:

A. Scaling rhythmically

B. Scaling max

C. Scaling out

D. Scaling up

**Suggested Answer:** *C*

*Community vote distribution*

C (94%)

**rakshithaK** Highly Voted 3 years, 3 months ago

Scale up by resizing a warehouse. Scale out by adding clusters to a multi-cluster warehouse (requires Snowflake Enterprise Edition or higher).

upvoted 8 times

**vitorsilva6** Highly Voted 2 years, 11 months ago

The questions are good, but the accepted answers are BAD.

upvoted 6 times

**Sudhansu21** Most Recent 1 year, 8 months ago

Scaling out refers to adding more nodes or clusters to a system to handle increased workload. In Snowflake, increasing the maximum number of clusters in a Multi-Cluster Warehouse is an example of scaling out because it adds more compute resources (clusters) to distribute the workload.

upvoted 1 times

**VLMBL_94** 2 years, 3 months ago

scaling out/in - multi cluster VWH ; scaling up/down : vwh resizing.

upvoted 1 times

**Tushar0807** 2 years, 3 months ago

Option C

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

Scale out is correct answer

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct answer:C

upvoted 1 times

**mnaik45** 2 years, 8 months ago

Scaling out is correct

upvoted 3 times

**_yyukta** 2 years, 8 months ago

C. Scaling out

upvoted 1 times

**Nilesh_Gore0310** 2 years, 11 months ago

I am seeing wrong answers for at least one question in every 10 questions.

upvoted 5 times

**cbji** 3 years, 3 months ago

Snowflake supports two ways to scale warehouses: Scale up by resizing a warehouse. Scale out by adding clusters to a multi-cluster warehouse (requires Snowflake Enterprise Edition or higher).

upvoted 3 times

**YATISH17** 3 years, 4 months ago

SCALLING OUT

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

D is correct

upvoted 1 times

**Harpreetwal** 3 years, 6 months ago

Scale out by adding more, and scale up is increasing hardware power for the single system

upvoted 3 times

**RajuNaik** 3 years, 6 months ago

Scaling out is correct

upvoted 2 times

**SteveCrown** 3 years, 6 months ago

Scaling Max Clusters. Stephen Okon Increasing the maximum number of clusters in a Multi-Cluster Warehouse is an example of: A. Scaling rhythmically B. Scaling max C. Scaling out D. Scaling up The correct answer is B. Scaling max. Scaling max refers to increasing the maximum limit of resources that can be allocated to a system or application. In the context of a Multi-Cluster Warehouse, increasing the maximum number of clusters would require more resources to be allocated to the system in order to accommodate the additional clusters. This is an example of scaling max because it involves increasing the maximum limit of resources that can be allocated to the Multi-Cluster Warehouse. Scaling rhythmically, on the other hand, refers to increasing or decreasing resources in a gradual and predictable manner based on the workload of the system. Scaling out involves adding more instances of a system or application to distribute the workload across multiple nodes. Scaling up involves increasing the resources of a single node in order to handle a larger workload.this completely make sense but couldn't find it in the documentation

upvoted 1 times

**SteveCrown** 3 years, 6 months ago

Scaling Max Clusters.

Stephen Okon Increasing the maximum number of clusters in a Multi-Cluster Warehouse is an example of: A. Scaling rhythmically B. Scaling max C. Scaling out D. Scaling up

The correct answer is B. Scaling max.

Scaling max refers to increasing the maximum limit of resources that can be allocated to a system or application. In the context of a Multi-Cluster Warehouse, increasing the maximum number of clusters would require more resources to be allocated to the system in order to accommodate the additional clusters. This is an example of scaling max because it involves increasing the maximum limit of resources that can be allocated to the Multi-Cluster Warehouse.

Scaling rhythmically, on the other hand, refers to increasing or decreasing resources in a gradual and predictable manner based on the workload of the system. Scaling out involves adding more instances of a system or application to distribute the workload across multiple nodes. Scaling up involves increasing the resources of a single node in order to handle a larger workload.

upvoted 1 times

**girgir** 3 years, 1 month ago

this completely make sense but couldn't find it in the documentation

upvoted 3 times

**Ayush24** 3 years, 6 months ago

Correct answer : C https://docs.snowflake.com/en/user-guide/warehouses-considerations#scaling-up-vs-scaling-out Snowflake supports two ways to scale warehouses: 1. Scale up by resizing a warehouse. 2. Scale out by adding clusters to a multi-cluster warehouse (requires Snowflake Enterprise Edition or higher).

upvoted 3 times

Question #42 *Topic 1*

Which statement best describes Snowflake tables?

A. Snowflake tables are logical representations of underlying physical data

B. Snowflake tables are the physical instantiation of data loaded into Snowflake

C. Snowflake tables require that clustering keys be defined to perform optimally

D. Snowflake tables are owned by a user

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Sudhansu21** Most Recent 1 year, 8 months ago

Snowflake, tables are logical structures that represent the data stored in the underlying cloud storage . Snowflake manages the physical storage and optimization of the data transparently.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A is Correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A is correct

upvoted 1 times

**Colourseun** 2 years, 11 months ago

The best statement that describes Snowflake tables is Snowflake tables are logical representations of underlying physical data. Snowflake tables do not require that clustering keys be defined to perform optimally. However, defining clustering keys can improve the performance of queries that filter or sort the data. Snowflake tables can be owned by a user or by a role. Therefore, the correct answer is A. Snowflake tables are logical representations of underlying physical data.

upvoted 4 times

**Gwo** 3 years, 1 month ago

C seems correct too. Snowflake tables require clustering keys to perform optimally. Clustering keys help manage micro partitions for a table, which in turn helps in query performance by means of partition pruning.Clustering is defined automatically by SF, not by the user, and clustering is an output of MPs, which are based on contiguous records. You can overwrite clustering keys but it is not recommended.

upvoted 1 times

**Gwo** 3 years, 1 month ago

C seems correct too. Snowflake tables require clustering keys to perform optimally. Clustering keys help manage micro partitions for a table, which in turn helps in query performance by means of partition pruning.

upvoted 1 times

**56d43d6** 1 year, 10 months ago

Clustering is defined automatically by SF, not by the user, and clustering is an output of MPs, which are based on contiguous records. You can overwrite clustering keys but it is not recommended.

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**alanhuangjob** 3 years, 7 months ago

A is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

A is correct

upvoted 2 times

**Percy2112** 4 years, 2 months ago

A is correct answer.

upvoted 2 times

**Rajokkiyam** 4 years, 9 months ago

Answer A

upvoted 3 times

Question #43 *Topic 1*

Which item in the Data Warehouse migration process does not apply in Snowflake?

A. Migrate Users

B. Migrate Schemas

C. Migrate Indexes

D. Build the Data Pipeline

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**lostri** Highly Voted 4 years, 9 months ago

Snowflake does not use indexes. Answer is C

upvoted 5 times

**Nittin** Most Recent 1 year, 11 months ago

it has micro partition and automatic clustering so doesnt use indexes

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

Migrate Indexes is correct answer

upvoted 1 times

**_yyukta** 2 years, 8 months ago

Migrate indexes

upvoted 1 times

**zephyrvz** 3 years ago

C is correct

upvoted 2 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

C is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

C is correct

upvoted 1 times

**Sundarr** 4 years, 9 months ago

Answer is C

upvoted 2 times

Question #44 *Topic 1*

Snowflake provides two mechanisms to reduce data storage costs for short-lived tables. These mechanisms are: (Choose two.)

A. Temporary Tables

B. Transient Tables

C. Provisional Tables

D. Permanent Tables

**Suggested Answer:** *A, B*

*Community vote distribution*

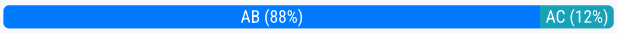


**jeny12138** Most Recent 1 year, 7 months ago

Temporary / provisional are true

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

Temporary and Transient Tables

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A and B

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

AB is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

AB is correct

upvoted 3 times

**Sundarr** 4 years, 9 months ago

A&B are correct

upvoted 2 times

**moon2351** 4 years, 10 months ago

Answer is correct.

upvoted 3 times

Question #45 *Topic 1*

What is the maximum compressed row size in Snowflake?

A. 8KB

B. 16MB

C. 50MB

D. 4000GB

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**RichMelbourne** Highly Voted 3 years, 10 months ago

Was on exam 20/12/2022

upvoted 7 times

**niho1993** Most Recent 1 year, 1 month ago

For columns of the following types, the maximum allowed length is 128 MB: VARCHAR VARIANT ARRAY OBJECT For columns of the following types, the maximum allowed length is 64 MB: BINARY GEOGRAPHY GEOMETRY

upvoted 1 times

**7b52dde** 1 year, 3 months ago

Historically, this limit for many of these "large object" (LOB) types was 16 MB (compressed). However, Snowflake has increased these limits in recent updates (BCR Bundles). Current (or recently increased) limits for individual column values: 128 MB for VARCHAR, VARIANT, ARRAY, and OBJECT. 64 MB for BINARY, GEOGRAPHY, and GEOMETRY.

upvoted 1 times

**jeny12138** 1 year, 7 months ago

16MB is correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

16 Mb is the max compressed size

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B.16MB

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**AtulRaje** 3 years, 9 months ago

https://docs.snowflake.com/en/user-guide/data-load-considerations-prepare.html#semi-structured-data-size-limitations 16MB per row captured in the variant field

upvoted 2 times

**Tapasgup007** 3 years, 11 months ago

B is correct

upvoted 1 times

**Sundarr** 4 years, 9 months ago

B is correct

upvoted 1 times

**arpit_dataguy** 4 years, 9 months ago

MAX size for compressed data for semi-structured data type is also 16 MB. That means if our table has a semi-structured data type column with 16MB data other columns will have NULL data.

upvoted 3 times

Question #46 *Topic 1*

Which of the following are main sections of the top navigation of the Snowflake Web Interface (UI)? (Choose three.)

A. Databases

B. Tables

C. Warehouses

D. Worksheets

**Suggested Answer:** *A, C, D*

*Community vote distribution*

ACD (100%)

**56d43d6** Highly Voted 1 year, 10 months ago

This is an old question now. Worksheets are now contained inside "Projects" and it isn't called the SF Web Interface any more (deprecated), it's now called SnowSight.And Tables are stored under Databases

upvoted 5 times

**56d43d6** 1 year, 10 months ago

And Tables are stored under Databases

upvoted 1 times

**Mallikharjuna452** 4 months ago

Variant

upvoted 2 times

**MikasSi** 1 year, 8 months ago

Deprecated question

upvoted 2 times

**Tapasgup007** 4 months ago

Selected Answer: D

D is correct

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

A,C,D are correct

upvoted 1 times

**hongy** Most Recent 4 months ago

Selected Answer: D

Answer D is correct.

upvoted 2 times

**_yyukta** 2 years, 8 months ago

ACD are correct

upvoted 1 times

**jeny12138** 4 months ago

Selected Answer: D

D is correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

ACD is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

ACD is correct

upvoted 2 times

**Sundarr** 4 years, 9 months ago

Answer is correct.

upvoted 2 times

**moon2351** 4 years, 10 months ago

Answer is correct.

upvoted 3 times

Question #47 *Topic 1*

What is the recommended Snowflake data type to store semi-structured data like JSON?

A. VARCHAR

B. RAW

C. LOB

D. VARIANT

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**jeny12138** Most Recent 1 year, 7 months ago

D is correct

upvoted 1 times

**Sudhansu21** 1 year, 8 months ago

D. VARIANT The VARIANT data type is specifically designed to efficiently store and query semi-structured data such as JSON, Avro, Parquet, and XML. It allows Snowflake to handle complex data structures and perform schema-on-read operations, making it ideal for storing and processing semi-structured data.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

Variant

upvoted 2 times

**_yyukta** 2 years, 8 months ago

D . Variant

upvoted 1 times

**Marge23** 3 years, 2 months ago

VARIANT

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**PoornaChandraIn** 3 years, 3 months ago

Object, Variant and Array are the semi structured data types supports in snowflake.

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

D is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

D is correct

upvoted 2 times

**hongy** 4 years, 7 months ago

Answer D is correct.

upvoted 2 times

**moon2351** 4 years, 10 months ago

Answer is correct.

upvoted 3 times

Question #48 *Topic 1*

Which of the following statements are true of Snowflake releases: (Choose two.)

A. They happen approximately weekly

B. They roll up and release approximately monthly, but customers can request early release application

C. During a release, new customer requests/queries/connections transparently move over to the newer version

D. A customer is assigned a 30 minute window (that can be moved anytime within a week) during which the system will be unavailable and customer is upgraded

**Suggested Answer:** *A, C*

*Community vote distribution*

AC (93%) (3

**PrithviPatil** Highly Voted 4 years, 10 months ago

https://docs.snowflake.com/en/user-guide/intro-releases.html

upvoted 13 times

**jeny12138** Most Recent 1 year, 7 months ago

ac is correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A and C

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct answer:AC

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

A. They happen approximately weekly C. During a release, new customer requests/queries/connections transparently move over to the newer version Snowflake operates on a continuous delivery model, with updates and improvements being rolled out frequently, typically on a weekly basis. These updates are designed to be seamless, with new requests, queries, and connections being transparently handled by the updated version without causing disruption to the users. Therefore, the correct options are: A. They happen approximately weekly C. During a release, new customer requests/queries/connections transparently move over to the newer version

upvoted 1 times

**Stevec80** 2 years, 5 months ago

https://docs.snowflake.com/en/user-guide/intro-releases.html We deploy new releases each week. The deployments happen transparently in the background; users experience no downtime or disruption of service, and are always assured of running on the most-recent release with access to the latest features.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A and C

upvoted 1 times

**samudu141** 3 years, 3 months ago

ac https://docs.snowflake.com/en/user-guide/intro-releases.html

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

CD is correct

upvoted 1 times

**c0d3g** 3 years, 8 months ago

Check the doc: https://docs.snowflake.com/en/user-guide/intro-releases

upvoted 2 times

**Raju039** 3 years, 8 months ago

AC is correct

upvoted 1 times

**KarBiswa** 3 years, 9 months ago

as per the documentation

upvoted 1 times

**ajay_1233456** 3 years, 9 months ago

AC are correct

upvoted 2 times

**BigDataBB** 3 years, 9 months ago

https://docs.snowflake.com/en/user-guide/intro-releases.html

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

AC are correct

upvoted 1 times

**MTAD** 4 years, 1 month ago

The answer looks A& C according to the documents on :- https://docs.snowflake.com/en/user-guide/intro-releases.html "....Snowflake is committed to providing a seamless, always up-to-date experience for our users while also delivering ever-increasing value through rapid development and continual innovation. To meet this commitment, we deploy new releases each week." "... This topic describes the process we follow for weekly releases, including the option to request 24-hour early access for Enterprise Edition (and higher) accounts to enable additional release testing (if desired)."

upvoted 1 times

**Diegovaz** 4 years, 4 months ago

Only C is correct, because it's not approximately weekly, it's weekly without errorI'd assume that that's what they mean with "approximately"...weekly pushes can be skipped if there if a full release scheduled nearby

upvoted 1 times

**Diegovaz** 4 years, 4 months ago

**Selected Answer: C**

Only C is correct, because it's not approximately weekly, it's weekly without error

upvoted 1 times

**oscarglob** 3 years, 1 month ago

I'd assume that that's what they mean with "approximately"...

upvoted 1 times

**DRAC25** 4 years ago

weekly pushes can be skipped if there if a full release scheduled nearby

upvoted 1 times

Question #49 *Topic 1*

Which of the following are common use cases for zero-copy cloning? (Choose three.)

A. Quick provisioning of Dev and Test/QA environments

B. Data backups

C. Point in time snapshots

D. Performance optimization

**Suggested Answer:** *A, B, C*

*Community vote distribution*

ABC (100%)

**jeny12138** Most Recent 1 year, 7 months ago

is not for optimization

upvoted 1 times

**someshVashistha** 1 year, 10 months ago

A,B,C are correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A,B,C are correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A B and C are correct

upvoted 2 times

**Honeybii** 3 years, 1 month ago

I think replication better represents data backup. Zero copy clone does not create physical backups, it's two logical data objects pointing to the same physical location.I think it's meant as a backup in the same way that it's a snapshot. It has data up to a certain point in time, then when the original table is modified, you still have the clone.

upvoted 3 times

**Honeybii** 3 years, 1 month ago

I think replication better represents data backup. Zero copy clone does not create physical backups, it's two logical data objects pointing to the same physical location.

upvoted 3 times

**oscarglob** 3 years, 1 month ago

I think it's meant as a backup in the same way that it's a snapshot. It has data up to a certain point in time, then when the original table is modified, you still have the clone.

upvoted 1 times

**samudu141** 3 years, 3 months ago

ABC correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

ABC is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

ABC is correct

upvoted 3 times

**MinuteMaid** 4 years, 7 months ago

ABC chocolate

upvoted 2 times

**Fosterchild123** 4 years, 8 months ago

Correct answer

upvoted 2 times

Question #50 *Topic 1*

If a Small Warehouse is made up of 2 servers/cluster, how many servers/cluster make up a Medium Warehouse?

A. 4

B. 16

C. 32

D. 128

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Fosterchild123** Highly Voted 4 years, 8 months ago

Correct answer. With each warehouse size incremental increase everything doubles including cost

upvoted 6 times

**jeny12138** Most Recent 1 year, 7 months ago

2*2=4 is correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

4 is correct Answer

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A is correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**SV1122** 3 years, 10 months ago

Was on exam Dec 31st, 2022

upvoted 4 times

**Tapasgup007** 3 years, 11 months ago

A is correct

upvoted 1 times

Question #51 *Topic 1*

True or False: When a data share is established between a Data Provider and a Data Consumer, the Data Consumer can extend that data share to other Data
Consumers.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**jeny12138** Most Recent 1 year, 7 months ago
consumers cannot share
upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago
B.False
upvoted 1 times

**p22_nilesh** 2 years, 5 months ago
Shared object cannot be further shared. Hence its FALSE
upvoted 2 times

**_yyukta** 2 years, 8 months ago
B. False
upvoted 1 times

**Srinivas9362** 2 years, 10 months ago
False, giving acces to perticular consumer is only more but can tht consumer can share that to another consumer
upvoted 1 times

**Marge23** 3 years, 2 months ago
FALSE General Limitations for Shared Databases¶ Shared databases have the following limitations for consumers: Shared databases are read-only. Users in a consumer account can view/query data, but cannot insert or update data, or create any objects in the database. The following actions are not supported: Creating a clone of a shared database or any schemas/tables in the database. Time Travel for a shared database or any schemas/tables in the database. Editing the comments for a shared database. Shared databases and all the objects in the database cannot be re-shared with other accounts. Shared databases cannot be replicated.
upvoted 4 times

**MultiCloudIronMan** 3 years, 3 months ago
Correct
upvoted 1 times

**Mayuri917233** 3 years, 5 months ago
B is correct
upvoted 1 times

**jw1806** 3 years, 7 months ago
yeah, you can not re-share your share. but in real live, I did it using dbt.
upvoted 2 times

**Tapasgup007** 3 years, 11 months ago
B is correct
upvoted 1 times

**Sundarr** 4 years, 9 months ago
Consumer can't re-shared DB object to another consumer.
upvoted 4 times

**Rajokkiyam** 4 years, 9 months ago

Answer False.

upvoted 2 times

**grabc** 4 years, 9 months ago

It makes sense not allowed to

upvoted 1 times

**Tapsinha** 4 years, 9 months ago

Consumer can't re-shared DB object to another consumer. https://docs.snowflake.com/en/user-guide/data-share-consumers.html

upvoted 1 times

Question #52 *Topic 1*

Which is true of Snowflake network policies? A Snowflake network policy: (Choose two.)

A. Is available to all Snowflake Editions

B. Is only available to customers with Business Critical Edition

C. Restricts or enables access to specific IP addresses

D. Is activated using an ג€ALTER DATABASEג€ command

**Suggested Answer:** *A, C*

*Community vote distribution*

AC (100%)

**jeny12138** Most Recent 1 year, 7 months ago

AC is correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

C and A are Correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A and C

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

AC are correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

AC are correct

upvoted 2 times

**Sundarr** 4 years, 9 months ago

Sample SQL: create or replace network policy mypolicy1 allowed_ip_list=('192.168.1.0/24') blocked_ip_list=('192.168.1.99');

upvoted 4 times

**moon2351** 4 years, 10 months ago

Answer is correct.

upvoted 3 times

Question #53 *Topic 1*

True or False: Snowflake charges additional fees to Data Providers for each Share they create.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (98%) 2

**jjordan** Highly Voted 4 years, 10 months ago

According to the documentation sited: "The costs for sharing data with Snowflake are minimal and straightforward. Data providers simply pay Snowflake for the data they store, and data consumers pay for the compute resources their queries consume."

upvoted 17 times

**Prajdeo** Highly Voted 3 years, 8 months ago

https://docs.snowflake.com/en/user-guide/data-sharing-intro According to docs, there is no separate cost to provider for creating the Share. Shared data does not take up any storage in a consumer account and therefore does not contribute to the consumer's monthly data storage charges. The only charges to consumers are for the compute resources (i.e. virtual warehouses) used to query the shared data. Because no data is copied or exchanged, Secure Data Sharing setup is quick and easy for providers and access to the shared data is near-instantaneous for consumers

upvoted 7 times

**jeny12138** Most Recent 1 year, 7 months ago

not charge additional fee

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B False

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct answer:B no additional cost other than storage

upvoted 1 times

**Chhatresh** 2 years, 6 months ago

This question is incomplete - it depends if its READER account (on cross cloud platform/non SF platform) or if its on same SF platform (same Cloud) consumer account. For Reader Account billing is on Provider account, for Consumer Account (same SF cloud) - its on Consumer Account

upvoted 4 times

**_yyukta** 2 years, 8 months ago

B. False

upvoted 1 times

**VASI250** 2 years, 10 months ago

FALSE. There is no additional cost based on number of shares created instead the charges are based on storage used.

upvoted 2 times

**ifmateus** 3 years, 2 months ago

https://community.snowflake.com/s/question/0D50Z00008MpYp8SAF/how-the-cost-for-data-storage-for-data-shared-is-calculated-in-snowflake-data-sharing

upvoted 1 times

**Amitsnowflake** 3 years, 2 months ago

"The costs for sharing data with Snowflake are minimal and straightforward. Data providers simply pay Snowflake for the data they store, and data consumers pay for the compute resources their queries consume."

upvoted 2 times

**Venkat_ithakota** 3 years, 3 months ago

No, there is no additional charge. This is because, when you share data, no actual data is copied, so there is no additional storage required and therefore no associated additional cost. Even the process of creating the share and granting privileges for databases and other supported database objects (schemas, UDFs, tables, and views) to a share does not incur any direct cost since these are metadata operations and are thus part of Cloud Services. Snowflake credits are used to pay for the usage of the cloud services that exceeds 10% of the daily usage of the compute resources. Usage for cloud services is charged only if the daily consumption of cloud services exceeds 10% of the daily usage of the compute resources. If data is shared outside the region where the Snowflake account is hosted then will be additional storage costs for replication.

upvoted 5 times

**Leinho** 3 years, 4 months ago

There is ne additional fee for datashare itself

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**piconcomeback** 3 years, 6 months ago

There is no extra cost in a share

upvoted 1 times

**c0d3g** 3 years, 8 months ago

No additional cost for share.

upvoted 1 times

**Josa2** 3 years, 8 months ago

As documentation there are no additional costs for the share

upvoted 2 times

**KarBiswa** 3 years, 9 months ago

Sharing is free as it within Snowflake accounts

upvoted 1 times

Question #54 *Topic 1*

Query results are stored in the Result Cache for how long after they are last accessed, assuming no data changes have occurred?

A. 1 Hour

B. 3 Hours

C. 12 hours

D. 24 hours

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**jeny12138** Most Recent 1 year, 7 months ago

D is correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

24 hours

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D is correct

upvoted 2 times

**Sarav13** 2 years, 8 months ago

24 hours

upvoted 2 times

**Marge23** 3 years, 2 months ago

24 HOURS

upvoted 2 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

D is correct

upvoted 2 times

**Tapasgup007** 3 years, 11 months ago

D is correct

upvoted 2 times

**moon2351** 4 years, 10 months ago

Answer is correct. "Which holds the results of every query executed in the past 24 hours. These are available across virtual warehouses, so query results returned to one user is available to any other user on the system who executes the same query, provided the underlying data has not changed."

upvoted 4 times

Question #55 *Topic 1*

A role is created and owns 2 tables. This role is then dropped. Who will now own the two tables?

A. The tables are now orphaned

B. The user that deleted the role

C. SYSADMIN

D. The assumed role that dropped the role

**Suggested Answer:** *D*

*Community vote distribution*

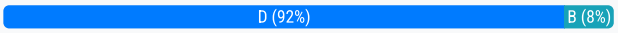


**uncle_benz** Highly Voted 3 years, 12 months ago

A user can't own a table. Ownership of any objects owned by the dropped role is transferred to the *role* that executes the DROP ROLE command.

upvoted 7 times

**jeny12138** Most Recent 1 year, 7 months ago

who drop table own this table

upvoted 1 times

**PrasantSadangi** 1 year, 7 months ago

Ownership of any objects owned by the dropped role is transferred to the role that executes the DROP ROLE command.

https://docs.snowflake.com/en/sql-reference/sql/drop-role#usage-notes

upvoted 2 times

**rocky_vits1** 1 year, 10 months ago

https://docs.snowflake.com/en/sql-reference/sql/drop-role#usage-notes

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

D is Correct

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

D. The assumed role that dropped the role In Snowflake, when a role is dropped, the ownership of objects owned by that role (such as tables) is transferred to the role that performed the drop operation. Therefore, the assumed role that dropped the role will now own the two tables.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D is correct

upvoted 2 times

**BobFar** 3 years ago

Correct answer is D Ownership of any objects owned by the dropped role is transferred to the role that executes the DROP ROLE command.

https://docs.snowflake.com/en/sql-reference/sql/drop-role#usage-notes

upvoted 2 times

**Marge23** 3 years, 2 months ago

D. Dropped roles cannot be recovered; they must be recreated. Ownership of any objects owned by the dropped role is transferred to the role that executes the DROP ROLE command. To transfer ownership of each of these objects to a different role, use GRANT OWNERSHIP … COPY CURRENT GRANTS.

upvoted 1 times

**magnasiun1** 3 years, 4 months ago

D is the RIGHT ANSWER

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

D is correct

upvoted 1 times

**Josa2** 3 years, 8 months ago

D is the RIGHT ANSWER https://docs.snowflake.com/en/sql-reference/sql/drop-role#usage-notes Ownership of any objects owned by the dropped role is transferred to the role that executes the DROP ROLE command. To transfer ownership of each of these objects to a different role, use GRANT OWNERSHIP … COPY CURRENT GRANTS.

upvoted 3 times

**AndroJS** 3 years, 10 months ago

Users don´t own objects, role does so answer is D

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

D is correct

upvoted 1 times

**seaun** 4 years ago

D is correct

upvoted 2 times

**greenlever** 4 years ago

B is correct

upvoted 1 times

**greenlever** 4 years ago

B Correct

upvoted 1 times

Question #56 *Topic 1*

Which of the following connectors are available in the Downloads section of the Snowflake Web Interface (UI)? (Choose two.)

A. SnowSQL

B. ODBC

C. R

D. HIVE

**Suggested Answer:** *A, B*

*Community vote distribution*

AB (100%)

**jeny12138** Most Recent 1 year, 7 months ago

ab is correct

upvoted 1 times

**PrasantSadangi** 1 year, 7 months ago

Snowflake CLI Snowflake CLI Download ODBC Driver ODBC Download Snowpark API Snowpark Client Download Drivers Drivers and Libraries Scala and Java connectors Drivers and Libraries SnowCD Drivers and Libraries Snowpark ML Drivers and Libraries

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

Snowsql and ODBC

upvoted 1 times

**rogerek077** 2 years, 7 months ago

I guess this question has no correct answer right now. 'Download' button is missing in current UI version (Snowsight): https://community.snowflake.com/s/article/How-to-download-snowflake-connectors-and-drivers-from-the-web-interface And now, for all new users, this seems to be the only UI available: https://community.snowflake.com/s/question/0D5Do00000gX8ovKAC/how-to-get-classic-console

upvoted 3 times

**_yyukta** 2 years, 8 months ago

A and B

upvoted 1 times

**Marge23** 3 years, 2 months ago

AB the right ans.

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**malexandersalazar** 3 years, 5 months ago

AB, but I think that ODBC is not a connector. It is a driver.

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

AB is correct

upvoted 1 times

**SV1122** 3 years, 10 months ago

AB - Snow SQL and ODBC https://docs.snowflake.com/en/user-guide/snowflake-client-repository.html

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

AB are correct

upvoted 1 times

**Sundarr** 4 years, 9 months ago

A&B are correct

upvoted 2 times

**moon2351** 4 years, 10 months ago

Answer are A & B

upvoted 2 times

Question #57 *Topic 1*

Which of the following DML commands isn't supported by Snowflake?

A. UPSERT

B. MERGE

C. UPDATE

D. TRUNCATE TABLE

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**arpit_dataguy** Highly Voted 4 years, 3 months ago

There is nothing called UPSERT in snowflake. MERGE = UPDATE+INSERT+DELETE. Also, unlike other databases TRUNCATE is considered DML in snowflake.

upvoted 5 times

**jeny12138** Most Recent 1 year, 7 months ago

there is no upsert

upvoted 1 times

**_yyukta** 2 years, 2 months ago

A.UPSERT

upvoted 1 times

**Marge23** 2 years, 8 months ago

A. correct

upvoted 1 times

**v1tjkee** 2 years, 8 months ago

TRUNCATE TABLE is DDL operation, as far as I know it drops tables and recreates it.

upvoted 2 times

**Mayuri917233** 2 years, 11 months ago

A is correct

upvoted 2 times

**Tapasgup007** 3 years, 5 months ago

A is correct

upvoted 1 times

**Ccastan1** 4 years ago

What is the correct answer if not UPSERT?

upvoted 1 times

**arpit_dataguy** 4 years, 3 months ago

There is nothing called UPSERT in snowflake. MERGE = UPDATE+INSERT.

upvoted 1 times

**moon2351** 4 years, 4 months ago

Commands for inserting, deleting, updating, and merging data in Snowflake tables: INSERT INSERT (multi-table) MERGE UPDATE DELETE TRUNCATE TABLE

upvoted 1 times

Question #58 *Topic 1*

Which of the following statements is true of zero-copy cloning?

A. Zero-copy clones increase storage costs as cloning the table requires storing its data twice

B. All zero-copy clone objects inherit the privileges of their original objects

C. Zero-copy cloning is licensed as an additional Snowflake feature

D. At the instance/instant a clone is created, all micro-partitions in the original table and the clone are fully shared

**Suggested Answer:** *D*

*Community vote distribution*

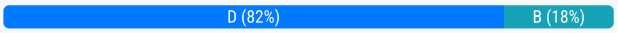


**MikeTQ** Highly Voted 3 years ago

Answer:D why not B? Note that the clone of the container itself (database or schema) does not inherit the privileges granted on the source container.

upvoted 5 times

**BobLamo** Most Recent 5 months ago

Access control privileges for cloned objects If the source object is a database or schema, the clone inherits all granted privileges on the clones of all child objects contained in the source object. D seems wrong because it looks 2-way, whereas clone lives its own life once created and the original schema/table won't track the clone.

upvoted 1 times

**jeny12138** 1 year, 7 months ago

no extra cost/ can't inherit privilege;

upvoted 2 times

**MohamedBen** 1 year, 11 months ago

The answer id B

upvoted 1 times

**lol632** 2 years, 11 months ago

It is D

upvoted 2 times

**AdahSimonis** 3 years, 1 month ago

Yes the answer is D

upvoted 1 times

**Marge23** 3 years, 2 months ago

Answer is D

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

D is the answer.

upvoted 1 times

**SV1122** 3 years, 10 months ago

Was on exam Dec 31st, 2022

upvoted 2 times

**DerekSlade** 3 years, 10 months ago

D is the answer.

upvoted 1 times

**aalonso** 3 years, 10 months ago

D is correct, but take into account that if the cloned or the original table make any change to the data, this change does not affect the other table. They are independent objects

upvoted 3 times

**Rajokkiyam** 4 years, 9 months ago

Answer D

upvoted 2 times

Question #59 *Topic 1*

True or False: When a user creates a role, they are initially assigned ownership of the role and they maintain ownership until it is transferred to another user.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*




**SV1122** Highly Voted 3 years, 10 months ago

By default, a newly-created role is not assigned to any user, nor granted to any other role. https://docs.snowflake.com/en/user-guide/security-access-control-overview.htmlThe question is regarding the assignment of ownership, not the assignment of the role. B is still correct though since the role is not initially owned by the user but rather by the role the user used during creation.

upvoted 9 times

**SV1122** 3 years, 10 months ago

Selected Answer: B

By default, a newly-created role is not assigned to any user, nor granted to any other role.

https://docs.snowflake.com/en/user-guide/security-access-control-overview.html

upvoted 9 times

**Badenhop** Highly Voted 3 years, 9 months ago

The question is regarding the assignment of ownership, not the assignment of the role. B is still correct though since the role is not initially owned by the user but rather by the role the user used during creation.

upvoted 19 times

**7b52dde** Most Recent 1 year, 3 months ago

it's a tricky question because "users" cannot own "roles". the role created is owned by the role that the user who created it was currently using. so role owned by role (assumed by user at the time)

upvoted 2 times

**jeny12138** 1 year, 7 months ago

when a user creates a role, they are not automatically assigned as the owner of that role

upvoted 1 times

**aguxkuro** 1 year, 7 months ago

In Snowflake, when a user creates a role, they are automatically assigned as the owner of that role

upvoted 1 times

**prakum** 1 year, 7 months ago

A. "OWNERSHIP is required to execute a CREATE OR ALTER ROLE statement for an existing role. OWNERSHIP is a special privilege on an object that is automatically granted to the role that created the object, but can also be transferred using the GRANT OWNERSHIP command to a different role by the owning role (or any role with the MANAGE GRANTS privilege). " https://docs.snowflake.com/en/sql-reference/sql/create-role

upvoted 1 times

**uairhahs** 1 year, 7 months ago

It is important to make the distinction between Users and Roles, users do not have any privelages including ownership this falls upon the role so rather than the user the ownership will go to the intial role used to create the new role.

upvoted 1 times

**akomal** 1 year, 7 months ago

yes, the creator maintains ownership until it is explicitly transferred to another role or user.

upvoted 1 times

**Siif97** 1 year, 9 months ago

When a user creates a role, they are initially assigned ownership to ensure that there is a clear and accountable owner who can manage the role's permissions and responsibilities. This ownership can be transferred to another user or role if needed, providing flexibility in managing access control.

upvoted 1 times

**MultiCloudIronMan** 1 year, 9 months ago

When a user creates a role, they are initially assigned ownership of the role and maintain ownership until it is transferred to another user. This aligns with the DAC model where the owner of an object has the discretion to grant access to other users or roles

upvoted 1 times

**volinets77** 1 year, 10 months ago

The ownership is defaulted to an actual role of the actual user

upvoted 1 times

**KIYREY** 1 year, 10 months ago

The role owner has full control

upvoted 1 times

**Mitchdu** 1 year, 10 months ago

Due to role hierarchy approach in snowflake, roles own other objects, not users. Ownership can be transferred to another role. So B.

upvoted 1 times

**MohamedBen** 1 year, 11 months ago

answer A

upvoted 1 times

**kentucky13caa** 1 year, 11 months ago

**Selected Answer: A**

The answer is A. True.

In Snowflake, when a user creates a role, they automatically become the owner of that role and retain ownership until it is explicitly transferred to another user. This is an important security concept in Snowflake's access control model:

Key points: - The creator of a role becomes its owner by default - Role ownership gives special privileges for managing that role - Only the role owner can grant or revoke role privileges (unless ownership is transferred) - Ownership transfer must be done explicitly using the GRANT OWNERSHIP command - Role ownership is distinct from role membership

This is important for the SnowPro Core exam as it relates to Snowflake's security model and role-based access control (RBAC) implementation.

upvoted 1 times

**kentucky13caa** 1 year, 11 months ago

The answer is A. True. In Snowflake, when a user creates a role, they automatically become the owner of that role and retain ownership until it is explicitly transferred to another user. This is an important security concept in Snowflake's access control model: Key points: - The creator of a role becomes its owner by default - Role ownership gives special privileges for managing that role - Only the role owner can grant or revoke role privileges (unless ownership is transferred) - Ownership transfer must be done explicitly using the GRANT OWNERSHIP command - Role ownership is distinct from role membership This is important for the SnowPro Core exam as it relates to Snowflake's security model and role-based access control (RBAC) implementation.Sounds convincing, but you need to provide evidence. Users do not create roles. Users are granted the ability to assume a role, and then a role creates another role. https://docs.snowflake.com/en/user-guide/security-access-control-overview#roles

upvoted 1 times

**56d43d6** 1 year, 10 months ago

Sounds convincing, but you need to provide evidence. Users do not create roles. Users are granted the ability to assume a role, and then a role creates another role. https://docs.snowflake.com/en/user-guide/security-access-control-overview#roles

upvoted 1 times

**khaled.mebhah** 1 year, 12 months ago

**True.** When a user creates a role in Snowflake, they are automatically assigned as the owner of that role. This ownership includes the ability to grant privileges, revoke privileges, and transfer ownership of the role. The user retains ownership of the role until it is explicitly transferred to another user, typically using the `GRANT OWNERSHIP` command. Ownership transfer is required if someone else needs to fully control or manage that role's permissions.

upvoted 1 times

**nk_24** 2 years ago

That's correct! In Snowflake, when a user creates a role, they are automatically assigned ownership of that role. This ownership includes the ability to grant and revoke privileges on the role, as well as the ability to transfer ownership to another user. The user retains ownership until they explicitly transfer it to someone else.

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct answer:B ownership cannot be transfered

upvoted 3 times

Question #60 *Topic 1*

The Query History in the Snowflake Web Interface (UI) is kept for approximately:

A. 60 minutes

B. 24 hours

C. 14 days

D. 30 days

E. 1 year

**Suggested Answer:** *C*

*Community vote distribution*

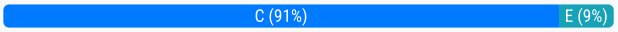


**LuciBhish** Highly Voted 4 years, 8 months ago

Answer C History page holds 14days data account usage->Access_history holds 1year data

upvoted 13 times

**Prajdeo** Highly Voted 3 years, 8 months ago

Answer should be C (14 days) As per docs - Query History page available under Activity Menu on Web UI, holds query data for 14 days. And you can Query_history view available in Account_Usage schema holds query data for 365 days https://docs.snowflake.com/en/user-guide/ui-snowsight-activity#query-history https://docs.snowflake.com/en/sql-reference/account-usage/query_history

upvoted 9 times

**jeny12138** Most Recent 1 year, 7 months ago

14 days is correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

14 Days

upvoted 1 times

**DaruB** 2 years, 4 months ago

The answer is C: https://docs.snowflake.com/en/user-guide/ui-snowsight-activity

upvoted 1 times

**pranalig** 2 years, 4 months ago

14 days

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

C. 14 days The Query History in the Snowflake Web Interface (UI) is kept for approximately 14 days. This allows users to review and analyze queries executed within that timeframe.

upvoted 1 times

**Stevec80** 2 years, 5 months ago

14 days via web interfaces and 1 year via SQL query

upvoted 1 times

**examtopics_strata** 3 years, 3 months ago

C correct

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

E is the answer

upvoted 1 times

**pedroramos** 3 years, 6 months ago

14 days

upvoted 1 times

**KarBiswa** 3 years, 9 months ago

correct for the Web UI perspective

upvoted 1 times

**rajivgupta780** 3 years, 9 months ago

https://docs.snowflake.com/en/user-guide/warehouses-overview.html

upvoted 1 times

**ajay_1233456** 3 years, 9 months ago

14 days

upvoted 1 times

**SV1122** 3 years, 10 months ago

This Account Usage view can be used to query Snowflake query history by various dimensions (time range, session, user, warehouse, etc.) within the last 365 days (1 year). https://docs.snowflake.com/en/sql-reference/account-usage/query_history.html

upvoted 2 times

**Tapasgup007** 3 years, 11 months ago

The History History tab page allows you to view and drill into the details of all queries executed in the last 14 days

upvoted 1 times

**stornati** 3 years, 11 months ago

QUERY_HISTORY View This Account Usage view can be used to query Snowflake query history by various dimensions (time range, session, user, warehouse, etc.) within the last 365 days (1 year). The view is available in both the ACCOUNT_USAGE and READER_ACCOUNT_USAGE schemas with the following differences:

upvoted 2 times

Question #61 *Topic 1*

To run a Multi-Cluster Warehouse in auto-scale mode, a user would:

A. Configure the Maximum Clusters setting to ɹ€Auto-Scaleɹ€

B. Set the Warehouse type to ɹ€Autoɹ€

C. Set the Minimum Clusters and Maximum Clusters settings to the same value

D. Set the Minimum Clusters and Maximum Clusters settings to the different values

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**moon2351** Highly Voted 4 years, 10 months ago

If you set the minimum cluster count less than the maximum cluster count, then the warehouse runs in Auto-scale mode.

upvoted 5 times

**jeny12138** Most Recent 1 year, 7 months ago

max and min are different vales

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

D is Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

D is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

D is correct

upvoted 1 times

**kjdfsafc87456** 4 years, 2 months ago

The answer is correct as per Auto-scale: This mode is enabled by specifying different values for maximum and minimum number of clusters. In this mode, Snowflake starts and stops clusters as needed to dynamically manage the load on the warehouse: As the number of concurrent user sessions and/or queries for the warehouse increases, and queries start to queue due to insufficient resources, Snowflake automatically starts additional clusters, up to the maximum number defined for the warehouse. Similarly, as the load on the warehouse decreases, Snowflake automatically shuts down clusters to reduce the number of running clusters and, correspondingly, the number of credits used by the warehouse. refer: https://docs.snowflake.com/en/user-guide/warehouses-multicluster.html

upvoted 1 times

Question #62 *Topic 1*

Which of the following terms best describes Snowflake's database architecture?

A. Columnar shared nothing

B. Shared disk

C. Multi-cluster, shared data

D. Cloud-native shared memory

**Suggested Answer:** *C*

*Community vote distribution*

C (94%)

**BigDataBB** Highly Voted 3 years, 11 months ago

Built from the ground up for the cloud, Snowflake's unique multi-cluster shared data architecture delivers the performance, scale, elasticity, and concurrency today's organizations require. https://www.snowflake.com/product/architecture/

upvoted 17 times

**jeny12138** Most Recent 1 year, 7 months ago

c is correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

C is correct

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct answer:C

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

C. Multi-cluster, shared data Snowflake's database architecture is best described as multi-cluster, shared data. This architecture separates compute and storage, allowing multiple compute clusters to access the same data concurrently and scale independently.

upvoted 2 times

**Yah020** 2 years, 6 months ago

It's answer C. yes you share your disk, but shared data is the best way to describe their architecture apparently. this is also a question in their Level Up: Snowflake Key Concepts & Architecture. Which ONE of the following terms BEST describes Snowflake's Architecture? Shared Disk Shared Nothing Shared Data Shared Memory Shared data was correct.

upvoted 1 times

**VenkataPhaniPavanKumar** 2 years, 7 months ago

Answer is B https://docs.snowflake.com/en/user-guide/intro-key-concepts you cannot find a word with multi-cluster in the entire page.

upvoted 1 times

**Arkcv** 3 years ago

Are these answers given by Exam topics or these answers are given by snowflake? Cause in second case its worrying what happens when they ask this question in exam

upvoted 2 times

**Rajivnb** 3 years, 4 months ago

It should be B. Its a shared disk architecture. For the multi cluster to access this data first the physcial disk sotring the data should be shareable. So B should be correct. Multicluster shared data comes in picture only when question is about data processing which in this case its a MPP.

upvoted 3 times

**magnasiun1** 3 years, 4 months ago

https://www.snowflake.com/product/architecture/?lang=pt-br Built from the ground up for the cloud, Snowflake's unique multi-cluster shared data architecture delivers the performance, scale, elasticity, and concurrency today's organizations require.

upvoted 1 times

**shujaz** 3 years, 6 months ago

B because question is about Snowflake's database architecture only

upvoted 1 times

**RajuNaik** 3 years, 6 months ago

C is correct

upvoted 1 times

**alanhuangjob** 3 years, 7 months ago

Answer should be C

upvoted 1 times

**Sujay316** 3 years, 9 months ago

Answer is B Snowflake's architecture is a hybrid of traditional shared-disk and shared-nothing database architectures. Similar to shared-disk architectures, Snowflake uses a central data repository for persisted data that is accessible from all compute nodes in the platform. But similar to shared-nothing architectures, Snowflake processes queries using MPP (massively parallel processing) compute clusters where each node in the cluster stores a portion of the entire data set locally. This approach offers the data management simplicity of a shared-disk architecture, but with the performance and scale-out benefits of a shared-nothing architecture.

upvoted 4 times

**ajay_1233456** 3 years, 9 months ago

C is correct

upvoted 1 times

**sonu10906997** 3 years, 9 months ago

C is correct

upvoted 1 times

**SV1122** 3 years, 10 months ago

Was on exam Dec 31st, 2022

upvoted 1 times

Question #63 *Topic 1*

Which of the following are options when creating a Virtual Warehouse? (Choose two.)

A. Auto-drop

B. Auto-resize

C. Auto-resume

D. Auto-suspend

**Suggested Answer:** *C, D*

*Community vote distribution*

CD (100%)

**jeny12138** Most Recent 1 year, 7 months ago

CD is correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

D Is correct

upvoted 1 times

**harshkoshta** 3 years, 4 months ago

its D only, i think C is not present now,

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

CD are correct

upvoted 1 times

**alanhuangjob** 3 years, 7 months ago

C & D are correct

upvoted 2 times

**Tapasgup007** 3 years, 11 months ago

C & D are correct

upvoted 2 times

**kjdfsafc87456** 4 years, 2 months ago

The answer is correct. When creating a warehouse, you could assign auto-suspend time from 5 minutes to never and you could tick the auto-resume function ( it is ticked by default). The function means that if you stock use the warehouse, the warehouse will be suspended by your assigned auto suspended time and auto-resume when you start to query

upvoted 2 times

**moon2351** 4 years, 10 months ago

Answer is correct.

upvoted 3 times

Question #64 *Topic 1*

A Virtual Warehouse's auto-suspend and auto-resume settings apply to:

A. The primary cluster in the Virtual Warehouse

B. The entire Virtual Warehouse

C. The database the Virtual Warehouse resides in

D. The queries currently being run by the Virtual Warehouse

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**jeny12138** Most Recent 1 year, 7 months ago

entire warehouse

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B.Apply on entire virtual Warehouse

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**kjdfsafc87456** 4 years, 2 months ago

The answer CORRECT and those settings are applied on the entire VM

upvoted 1 times

**Percy2112** 4 years, 2 months ago

Appy on entire Virtual Warehouse.

upvoted 1 times

**MinuteMaid** 4 years, 7 months ago

The entire virtual warehouse is the correct answer ey

upvoted 2 times

Question #65 *Topic 1*

Fail-safe is unavailable on which table types? (Choose two.)

A. Temporary

B. Transient

C. Provisional

D. Permanent

**Suggested Answer:** *A, B*

*Community vote distribution*

AB (100%)

**jeny12138** Most Recent 1 year, 7 months ago

ab is correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A and B are correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

AB are correct

upvoted 1 times

**alanhuangjob** 3 years, 7 months ago

A&B are correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

AB are correct

upvoted 1 times

**kjdfsafc87456** 4 years, 2 months ago

Yes, the answer is AB as a permanent table needs high reliability while the first two are temporaries and the C is not the function of snflk.

upvoted 2 times

**DK69** 4 years, 5 months ago

AB is correct

upvoted 2 times

**moon2351** 4 years, 10 months ago

Transient and temporary tables have no Fail-safe period. As a result, no additional data storage charges are incurred beyond the Time Travel retention period.

upvoted 4 times

Question #66 *Topic 1*

Which of the following objects is not covered by Time Travel?

A. Tables

B. Schemas

C. Databases

D. Stages

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**jeny12138** Most Recent 1 year, 7 months ago

stage is not covered

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

Stages

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

D is correct

upvoted 1 times

**EmiB** 3 years, 7 months ago

D - STAGE

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

D is correct

upvoted 1 times

**DK69** 4 years, 5 months ago

D is correct

upvoted 2 times

**moon2351** 4 years, 10 months ago

Answer is correct. Restoring data-related objects (tables, schemas, and databases) that might have been accidentally or intentionally deleted.

upvoted 3 times

Question #67 *Topic 1*

True or False: Micro-partition metadata enables some operations to be completed without requiring Compute.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**vivek_nitd** Highly Voted 4 years, 8 months ago

A. like count(*), min or max etc

upvoted 13 times

**visionary4ever** Most Recent 1 year, 2 months ago

like estimation functions

upvoted 1 times

**jeny12138** 1 year, 7 months ago

it is correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A is corect

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 2 times

**Tapasgup007** 3 years, 11 months ago

A is correct

upvoted 1 times

**DK69** 4 years, 5 months ago

A is correct

upvoted 2 times

**moon2351** 4 years, 10 months ago

Answer is correct.

upvoted 2 times

Question #68 *Topic 1*

Which of the following commands are not blocking operations? (Choose two.)

A. UPDATE

B. INSERT

C. MERGE

D. COPY

**Suggested Answer:** *B, D*

*Community vote distribution*

BD (100%)

**AGuru2020** Highly Voted 4 years, 2 months ago

BD The following guidelines apply in most situations: COMMIT operations (including both AUTOCOMMIT and explicit COMMIT) lock resources, but usually only briefly. UPDATE, DELETE, and MERGE statements hold locks that generally prevent them from running in parallel with other UPDATE, DELETE, and MERGE statements. Most INSERT and COPY statements write only new partitions. Those statements often can run in parallel with other INSERT and COPY operations, and sometimes can run in parallel with an UPDATE, DELETE, or MERGE statement.

upvoted 18 times

**Tapsinha** Highly Voted 4 years, 3 months ago

Sorry INSERT and COPY. https://docs.snowflake.com/en/sql-reference/transactions.html Transactional operations acquire locks on a resource, such as a table, while that resource is being modified. Locks block other statements from modifying the resource until the lock is released. The following guidelines apply in most situations: COMMIT operations (including both AUTOCOMMIT and explicit COMMIT) lock resources, but usually only briefly. UPDATE, DELETE, and MERGE statements hold locks that generally prevent them from running in parallel with other UPDATE, DELETE, and MERGE statements. Most INSERT and COPY statements write only new partitions. Those statements often can run in parallel with other INSERT and COPY operations, and sometimes can run in parallel with an UPDATE, DELETE, or MERGE statement.

upvoted 7 times

**Alex_Ovidiu** Most Recent 2 years, 1 month ago

Here's why the other options are not typically blocking: B. INSERT (Non-blocking): Inserting new data typically doesn't block other operations on the same table. Snowflake manages insertions efficiently and allows concurrent access for reads and other inserts. D. COPY (Non-blocking in most cases): The COPY command for loading data is usually non-blocking. Snowflake manages the loading process in the background and allows other operations to proceed without interruption.

upvoted 2 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 2 times

**PoornaChandraIn** 2 years, 9 months ago

Answer is BD Explanation: Only UPDATE and DELETE DML operations are blocked when a table is locked. INSERT operations are NOT blocked. UPDATE, DELETE, and MERGE statements hold locks that generally prevent them from running in parallel with other UPDATE, DELETE, and MERGE statements. Most INSERT and COPY statements write only new partitions.

upvoted 2 times

**Mayuri917233** 2 years, 11 months ago

BD is correct

upvoted 2 times

**Tapasgup007** 3 years, 5 months ago

BD is correct

upvoted 2 times

**DK69** 3 years, 11 months ago

BD is correct

upvoted 2 times

**Tapsinha** 4 years, 3 months ago

INSERT MERGE https://docs.snowflake.com/en/sql-reference/transactions.html

upvoted 1 times

**Tapsinha** 4 years, 3 months ago

INSERT MERGE https://docs.snowflake.com/en/sql-reference/transactions.htmlWrong. In link you shared: ...Most INSERT and COPY statements write only new partitions. Those statements often can run in parallel with other INSERT and COPY operations... So I guess INSERT & COPY are correct answers.

upvoted 1 times

**brian_69420** 3 years, 1 month ago

Wrong. In link you shared: ...Most INSERT and COPY statements write only new partitions. Those statements often can run in parallel with other INSERT and COPY operations... So I guess INSERT & COPY are correct answers.

upvoted 1 times

Question #69 *Topic 1*

Which of the following is true of Snowpipe via REST API? (Choose two.)

A. You can only use it on Internal Stages

B. All COPY INTO options are available during pipe creation

C. Snowflake automatically manages the compute required to execute the Pipe's COPY INTO commands

D. Snowpipe keeps track of which files it has loaded

**Suggested Answer:** *C, D*

*Community vote distribution*

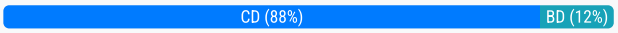


**redsky0** Highly Voted 4 years, 10 months ago

Answer is C & D. Not B as: https://docs.snowflake.com/en/sql-reference/sql/create-pipe.html#usage-notes All COPY INTO copy options are supported except for the following: FILES = ( 'file_name1' [ , 'file_name2', ... ] ) ON_ERROR = ABORT_STATEMENT SIZE_LIMIT = num PURGE = TRUE | FALSE (i.e. automatic purging while loading) FORCE = TRUE | FALSE

upvoted 38 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

C and D are correct

upvoted 1 times

**372110c** 2 years, 4 months ago

Correct Answer is CD

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct answer:CD since pipe is using serverless functionality

upvoted 2 times

**PrashantGupta1616** 2 years, 5 months ago

The correct options for what is true of Snowpipe via REST API are: C. Snowflake automatically manages the compute required to execute the Pipe's COPY INTO commands D. Snowpipe keeps track of which files it has loaded

upvoted 2 times

**Alex_Ovidiu** 2 years, 7 months ago

Here's why the other options are not accurate: A. You can only use it on Internal Stages: You can use Snowpipe via REST API with both Internal Stages (located within Snowflake storage) and External Stages (referencing data in external cloud storage providers like S3 or Azure Blob Storage). B. All COPY INTO options are available during pipe creation: Not all COPY INTO options are available through the Snowpipe REST API for pipe creation. While the core functionality is supported, some advanced options might be limited through the API compared to directly using the COPY INTO command within Snowflake.

upvoted 2 times

**oscarglob** 3 years ago

C & D: Snowpipe uses Snowflake-provided compute resources: https://docs.snowflake.com/user-guide/data-load-snowpipe-intro#calling-snowpipe-rest-endpoints

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

BD are correct

upvoted 2 times

**RajuNaik** 3 years, 6 months ago

Answer is C & D.

upvoted 1 times

**Prajdeo** 3 years, 8 months ago

Ans should be C & D

upvoted 1 times

**BigDataBB** 3 years, 9 months ago

CD are correct answer

upvoted 2 times

**Tapasgup007** 3 years, 11 months ago

CD are correct

upvoted 1 times

**bharathmv** 4 years, 1 month ago

C&D is right

upvoted 1 times

**DK69** 4 years, 5 months ago

CD are correct

upvoted 1 times

**AGuru2020** 4 years, 8 months ago

CD are correct answer

upvoted 2 times

**vivek_nitd** 4 years, 8 months ago

Snowpipe copy into has certain restrictions.

upvoted 2 times

**Rajokkiyam** 4 years, 9 months ago

Answer C D

upvoted 1 times

Question #70 *Topic 1*

Snowflake recommends, as a minimum, that all users with the following role(s) should be enrolled in Multi-Factor Authentication (MFA):

A. SECURITYADMIN, ACCOUNTADMIN, PUBLIC, SYSADMIN

B. SECURITYADMIN, ACCOUNTADMIN, SYSADMIN

C. SECURITYADMIN, ACCOUNTADMIN

D. ACCOUNTADMIN

**Suggested Answer:** *D*

*Community vote distribution*

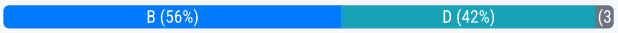


**moon2351** Highly Voted 4 years, 10 months ago

At a minimum, Snowflake strongly recommends that all users with the ACCOUNTADMIN role be required to use MFA.

upvoted 8 times

**newGodking** Most Recent 4 months, 2 weeks ago

I checked the official Snowflake docs. Snowflake official wording says MFA recommendation applies particularly to users with ACCOUNTADMIN, and can also be expanded to include users with SECURITYADMIN and SYSADMIN roles. B

upvoted 1 times

**8a28c1a** 11 months ago

https://docs.snowflake.com/en/user-guide/security-mfa In these cases, Snowflake recommends that each administrator set up a TOTP MFA method so that they can use an authenticator app as their second factor

upvoted 2 times

**d76ab49** 1 year, 1 month ago

https://docs.snowflake.com/en/user-guide/security-mfa In these cases, Snowflake recommends that each administrator set up a TOTP MFA method so that they can use an authenticator app as their second factor

upvoted 3 times

**HagarTheHorrible** 1 year, 4 months ago

see the red https://docs.snowflake.com/en/user-guide/security-mfa

upvoted 1 times

**dfd8f9f** 1 year, 6 months ago

this is the right answer.

upvoted 1 times

**BrunoRYamada** 1 year, 7 months ago

MINIMUM is the keyword

upvoted 1 times

**PrasantSadangi** 1 year, 7 months ago

If the account was created after the 2024_08 bundle was enabled, then all human users must enroll in MFA by default. For more information about this default behavior for new accounts, see the following:

upvoted 2 times

**PrasantSadangi** 1 year, 7 months ago

https://docs.snowflake.com/en/user-guide/security-mfa Attention At a minimum, Snowflake strongly recommends that all users with the ACCOUNTADMIN role be required to use MFA.

upvoted 2 times

**PrasantSadangi** 1 year, 7 months ago

https://www.snowflake.com/en/blog/multi-factor-identification-default/

upvoted 1 times

**tpiscaglia** 1 year, 7 months ago

Answer is B but recently Snowflake announced that they are pushing for ALL users to have MFA and are requiring it by default.
https://www.snowflake.com/en/blog/multi-factor-identification-default/

upvoted 3 times

**akomal** 1 year, 7 months ago

Snowflake recommends that, at a minimum, all users with the following roles should be enrolled in Multi-Factor Authentication (MFA): ACCOUNTADMIN SYSADMIN SECURITYADMIN

upvoted 2 times

**cgm_e** 1 year, 9 months ago

Snowflake Article, June 1 2024 We strongly recommend that customers enable multi-factor authentication (MFA) for admin users and users with access to sensitive data. This article provides instructions for reviewing which user accounts not currently leveraging MFA. Snowflake supports multi-factor authentication (i.e., MFA) to provide increased login security for users connecting to Snowflake. MFA support is provided as an integrated Snowflake feature, powered by the Duo Security service, which is managed completely by Snowflake. At a minimum, Snowflake strongly recommends that all users with the following system-defined roles enable MFA: ACCOUNTADMIN SECURITYADMIN SYSADMIN

upvoted 3 times

**MohamedBen** 1 year, 11 months ago

answer B

upvoted 1 times

**Shipra123** 1 year, 11 months ago

B. SECURITYADMIN, ACCOUNTADMIN, SYSADMIN as of June 1, 2024

upvoted 1 times

**jiriz** 2 years, 1 month ago

It doesn't make sense to me why they don't recommend securityadmin and sysadmin with MFA, but they really only recommend AccountAdmin. So D is correct.

upvoted 2 times

**ABRocks** 2 years, 2 months ago

B. SECURITYADMIN, ACCOUNTADMIN, SYSADMIN as of June 1, 2024

upvoted 4 times

**gany_man** 2 years, 1 month ago

SECURITYADMIN, ACCOUNTADMIN, SYSADMIN

upvoted 1 times

Question #71 *Topic 1*

When can a Virtual Warehouse start running queries?

A. 12am-5am

B. Only during administrator defined time slots

C. When its provisioning is complete

D. After replication

**Suggested Answer:** *C*

*Community vote distribution*

C (95%) (5%

**jjordan** Highly Voted 4 years, 10 months ago

https://docs.snowflake.com/en/user-guide/warehouses-overview.html "Warehouses can be started and stopped at any time."

upvoted 12 times

**pranalig** Most Recent 2 years, 4 months ago

Correct answer:C

upvoted 1 times

**Afzy** 3 years ago

Answer is C, when the synchronisation is complete

upvoted 1 times

**Sumannc** 3 years, 1 month ago

its C as the warehouse start executing the query after the provisioning. there is not specific time slot and no provision when an accountadmin specify a time window as long as the warehouse is running and the user have right privilege

upvoted 3 times

**ifmateus** 3 years, 2 months ago

The answer is C. "Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer." And "If cost and access are not an issue, enable auto-resume to ensure that the warehouse starts whenever needed. Keep in mind that there might be a short delay in the resumption of the warehouse due to provisioning." https://docs.snowflake.com/en/user-guide/warehouses-considerations

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

Selected Answer: B

B is correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correctincorrect

upvoted 1 times

**bbrandhh** 3 years, 3 months ago

incorrect

upvoted 3 times

**Raju039** 3 years, 8 months ago

C is correct

upvoted 1 times

**KarBiswa** 3 years, 9 months ago

Immediately starts after the warehouse resource has been alocated

upvoted 1 times

**ajay_1233456** 3 years, 9 months ago

C is correct

upvoted 1 times

**SV1122** 3 years, 10 months ago

Was on exam Dec 31st, 2022

upvoted 2 times

**baig123** 3 years, 10 months ago

C is correct, virtual warehouses can be configure to auto_resume=true/false, accordingly once it provision it start executing the queries.

upvoted 2 times

**Tapasgup007** 3 years, 11 months ago

C is correct

upvoted 1 times

**Graghu** 4 years, 6 months ago

Answer - C

upvoted 1 times

**yayaro** 4 years, 8 months ago

C is the answer

upvoted 2 times

**Sid460545** 4 years, 9 months ago

Answer is C

upvoted 2 times

Question #72 *Topic 1*

True or False: Users are able to see the result sets of queries executed by other users that share their same role.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*




**rsmurugan** Highly Voted 4 years, 9 months ago

Answer is B Viewing Query Results A user cannot view the result set from a query that another user executed. This behavior is intentional. For security reasons, only the user who executed a query can access the query results.

upvoted 11 times

**newGodking** Most Recent 4 months, 2 weeks ago

Users can see result sets from queries executed by other users when they are using the same role, because Snowflake result access is role-based in this case. Official Snowflake docs for RESULT_SCAN say that if a user is operating with the same role, they can use RESULT_SCAN to access query results. Snowflake stores query results for 24 hours.

upvoted 1 times

**goncalomsantos** 1 year, 8 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-considerations#viewing-query-results

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B False

upvoted 1 times

**dpexam2023** 3 years, 1 month ago

A user cannot view the result set from a query that another user executed. This behavior is intentional. For security reasons, only the user who executed a query can access the query results.

upvoted 1 times

**Marge23** 3 years, 2 months ago

B. FALSE

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**alfredofmt** 3 years, 7 months ago

A user cannot view the result set from a query that another user executed. This behavior is intentional. For security reasons, only the user who executed a query can access the query results. This behavior is not connected to the Snowflake access control model for objects. Even a user with the ACCOUNTADMIN role cannot view the results for a query run by another user. A user cannot view the result set from a query that another user executed, regardless of their assigned roles or privileges. Only the user who executed a query can access the query results.
https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html#viewing-query-results

upvoted 1 times

**ajay_1233456** 3 years, 9 months ago

A user cannot view the result set from a query that another user executed. This behavior is intentional. For security reasons, only the user who executed a query can access the query results.

upvoted 2 times

**BigDataBB** 3 years, 9 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html#viewing-query-results

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

B is correct

upvoted 1 times

Question #73 *Topic 1*

True or False: The user has to specify which cluster a query will run on in a multi-cluster Warehouse.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**jjordan** Highly Voted 4 years, 10 months ago

A user has to select a WAREHOUSE not a specific cluster within the warehouse.

upvoted 13 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B.false

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**Antz83** 4 years, 3 months ago

False - when select a warehouse you don't determine or have visibility to which cluster

upvoted 4 times

**aqmdla2002** 4 years, 3 months ago

Select the warehouse, not the cluster

upvoted 2 times

**yayaro** 4 years, 8 months ago

B is the answer

upvoted 1 times

**Sid460545** 4 years, 9 months ago

It is False.. A user doesn't need to specify warehouse

upvoted 1 times

**Sid460545** 4 years, 9 months ago

It is False.. A user doesn't need to specify warehouseUser can use warehouse, but doesn't need to specify cluster.. So False

upvoted 1 times

**Sid460545** 4 years, 9 months ago

User can use warehouse, but doesn't need to specify cluster.. So False

upvoted 1 times

**rsmurugan** 4 years, 9 months ago

Answer: B

upvoted 1 times

**Fab33** 4 years, 9 months ago

Selected Answer: B

upvoted 1 times

**PrithviPatil** 4 years, 10 months ago

A user has to select only the VW, not a specific cluster within the VW.

upvoted 3 times

Question #74 *Topic 1*

True or False: Pipes can be suspended and resumed.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*

A (88%) B (12%)

**jjordan** Highly Voted 4 years, 10 months ago

https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html#pausing-or-resuming-pipes Pausing or Resuming Pipes In addition to the pipe owner, a role that has the following minimum permissions can pause or resume the pipe:

upvoted 20 times

**ABRocks** Most Recent 2 years, 2 months ago

Pipe can be paused and resumed.

upvoted 1 times

**VLMBL_94** 2 years, 3 months ago

State of pipe can be suspend/resume

upvoted 1 times

**hotflake** 2 years, 3 months ago

WHAT IS A PIPE? ● Schema-level object ● Named definition of a COPY INTO ● Loads data into tables continuously from an ingestion queue ● Can be suspended/resumed, return status

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

YES, A IS CORRECT

upvoted 1 times

**LuYan18** 2 years, 6 months ago

Pausing or Resuming Pipes

upvoted 2 times

**Avi0407** 3 years ago

CAN PAUSE, BUT PIPE SUSPENDED...? I DIDN'T SEE ANYWHERE

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**nencheppan** 3 years, 7 months ago

A IS COORECT

upvoted 1 times

**Raju039** 3 years, 8 months ago

A is correct.

upvoted 1 times

**Jtravismiller** 3 years, 9 months ago

I lean towards B as there might be a difference between "suspended" and "paused".

upvoted 3 times

**ajay_1233456** 3 years, 9 months ago

A is correct

upvoted 1 times

**baig123** 3 years, 10 months ago

A is correct Answer.

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

A is correct

upvoted 1 times

**kjdfsafc87456** 4 years, 2 months ago

The answer is correct as the pipe can be stoped and resumed not b!

upvoted 1 times

**Sid460545** 4 years, 9 months ago

Can be suspended and resumed

upvoted 1 times

**rsmurugan** 4 years, 9 months ago

Answer A

upvoted 1 times

Question #75 *Topic 1*

Which of the following languages can be used to implement Snowflake User Defined Functions (UDFs)? (Choose two.)

A. Java

B. Javascript

C. SQL

D. Python

**Suggested Answer:** *B, C*

*Community vote distribution*

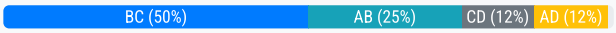


**jveag** Highly Voted 4 years, 1 month ago

ALL is correct. https://docs.snowflake.com/en/sql-reference/user-defined-functions.html

Snowflake currently supports the following languages for writing UDFs:

Java: A Java UDF lets you use the Java programming language to manipulate data and return either scalar or tabular results.

JavaScript: A JavaScript UDF lets you use the JavaScript programming language to manipulate data and return either scalar or tabular results.

Python: A Python UDF lets you use the Python programming language to manipulate data and return either scalar or tabular results.

SQL: A SQL UDF evaluates an arbitrary SQL expression and returns either scalar or tabular results.

upvoted 57 times

**jveag** Highly Voted 4 years, 1 month ago

ALL is correct. https://docs.snowflake.com/en/sql-reference/user-defined-functions.html Snowflake currently supports the following languages for writing UDFs: Java: A Java UDF lets you use the Java programming language to manipulate data and return either scalar or tabular results. JavaScript: A JavaScript UDF lets you use the JavaScript programming language to manipulate data and return either scalar or tabular results. Python: A Python UDF lets you use the Python programming language to manipulate data and return either scalar or tabular results. SQL: A SQL UDF evaluates an arbitrary SQL expression and returns either scalar or tabular results.https://docs.snowflake.com/en/sql-reference/udf-overview.html all are correctAll options are correct according to snowflake doc https://docs.snowflake.com/en/developer-guide/udf/udf-overview#supported-languages

upvoted 57 times

**KarBiswa** 3 years, 9 months ago

https://docs.snowflake.com/en/sql-reference/udf-overview.html all are correct

upvoted 6 times

**VenkataPhaniPavanKumar** 2 years, 9 months ago

All options are correct according to snowflake doc https://docs.snowflake.com/en/developer-guide/udf/udf-overview#supported-languages

upvoted 1 times

**moon2351** 4 years, 10 months ago

User-defined functions (UDFs) let you extend the system to perform operations that are not available through the built-in, system-defined functions provided by Snowflake. Snowflake currently supports the following languages for writing UDFs: SQL: A SQL UDF evaluates an arbitrary SQL expression and returns either scalar or tabular results. JavaScript: A JavaScript UDF lets you use the JavaScript programming language to manipulate data and return either scalar or tabular results. Java: A Java UDF lets you use the Java programming language to manipulate data and return either scalar or tabular results.

upvoted 8 times

**visionary4ever** Most Recent 1 year, 2 months ago

Snowflake now supports SQL, JavaScript, Java, Python

upvoted 1 times

**francisss** 1 year, 8 months ago

Only B, C is correct! https://docs.snowflake.com/en/developer-guide/udf/udf-overview --> Actually, Snowflake officially supports UDF implementations in four languages: (Java, JavaScript, SQL, Python) BUT, You must clearly understand the context in which this question is asked. Typically, Snowflake's certification exams and courses require JavaScript and SQL as default answers for the following reasons: ① Default Support : JavaScript and SQL UDFs are supported natively in all Snowflake editions without any additional setup or package installation. ② Java and Python UDFs have separate requirements : To use Java and Python UDFs, you need to activate the Snowpark feature provided by Snowflake, install related packages, and configure a separate environment. (https://docs.snowflake.com/en/developer-guide/udf/udf-overview)

upvoted 2 times

**MultiCloudIronMan** 1 year, 9 months ago

All correct ABCD

upvoted 1 times

**Yash_Chavan** 1 year, 9 months ago

Ans is all. But only SQL and JavaScript UDFs are sharable.

upvoted 1 times

**Srinivas_Vadakattu** 1 year, 11 months ago

Only Javascript and SQL

upvoted 2 times

**kentucky13caa** 1 year, 11 months ago

For the SnowPro Core exam, the two correct languages that can be used to implement Snowflake User Defined Functions (UDFs) are: B. Javascript C. SQL Here's the explanation: Supported languages for UDFs: 1. SQL: - Most common type of UDF - Written in SQL/ANSI syntax - Used for scalar and table functions - Runs natively in Snowflake 2. Javascript: - Supports more complex logic - Runs in secure sandbox environment - Can handle advanced calculations - Allows for more programming flexibility The incorrect options: - A. Java: Not supported for UDFs in Snowflake - D. Python: Not supported for UDFs (Note: Python is supported for Stored Procedures, but not for UDFs) Key points for the exam: - Only SQL and Javascript are supported for UDFs - Don't confuse UDF language support with Stored Procedure language support - JavaScript UDFs run in a secure execution environment - SQL UDFs are most commonly usedYour assertions are not supported by the documentation: https://docs.snowflake.com/en/developer-guide/udf/udf-overview#supported-languages UDFs support all the languages in this question. The question is probably old.

upvoted 5 times

**kentucky13caa** 1 year, 11 months ago

Selected Answer: BC

For the SnowPro Core exam, the two correct languages that can be used to implement Snowflake User Defined Functions (UDFs) are:

B. Javascript C. SQL

Here's the explanation:

Supported languages for UDFs: 1. SQL: - Most common type of UDF - Written in SQL/ANSI syntax - Used for scalar and table functions - Runs natively in Snowflake

2. Javascript: - Supports more complex logic - Runs in secure sandbox environment - Can handle advanced calculations - Allows for more programming flexibility

The incorrect options: - A. Java: Not supported for UDFs in Snowflake - D. Python: Not supported for UDFs (Note: Python is supported for Stored Procedures, but not for UDFs)

Key points for the exam: - Only SQL and Javascript are supported for UDFs - Don't confuse UDF language support with Stored Procedure language support - JavaScript UDFs run in a secure execution environment - SQL UDFs are most commonly used

upvoted 5 times

**56d43d6** 1 year, 10 months ago

Your assertions are not supported by the documentation: https://docs.snowflake.com/en/developer-guide/udf/udf-overview#supported-languages UDFs support all the languages in this question. The question is probably old.

upvoted 1 times

**Wydad_19** 1 year, 12 months ago

Issue here is that I cannot select all the options even though all answers are correct

upvoted 1 times

**dSilu** 2 years ago

All are correct. However, you can only share SQL and JavaScript UDFs

upvoted 1 times

**mackel** 2 years ago

All are correct https://docs.snowflake.com/en/developer-guide/udf/udf-overview#label-udf-supported-languages

upvoted 1 times

**hotflake** 2 years, 3 months ago

UDFs Can be written in o SQL o Python o Scala o Java o JavaScript

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

All options are correct

upvoted 1 times

**levisn1** 2 years, 4 months ago

All correct, capisce?

upvoted 1 times

**pranalig** 2 years, 4 months ago

correct answer:ABCD

upvoted 2 times

**Stevec80** 2 years, 5 months ago

All answers are correct. Also Java and Javascript are supported. https://docs.snowflake.com/en/developer-guide/udf/udf-overview#label-udf-supported-languages

upvoted 1 times

**sushant02** 2 years, 6 months ago

All four are correct.

upvoted 1 times

Question #76 *Topic 1*

When should you consider disabling auto-suspend for a Virtual Warehouse? (Choose two.)

A. When users will be using compute at different times throughout a 24/7 period

B. When managing a steady workload

C. When the compute must be available with no delay or lag time

D. When you do not want to have to manually turn on the Warehouse each time a user needs it

**Suggested Answer:** *B, C*

*Community vote distribution*

BC (100%)

**arpit_dataguy** Highly Voted 4 years, 9 months ago

From snowflake documentation : You might want to consider disabling auto-suspend for a warehouse if: You have a heavy, steady workload for the warehouse. You require the warehouse to be available with no delay or lag time. Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer.Source: https://docs.snowflake.com/en/user-guide/warehouses-considerations#automating-warehouse-suspension

upvoted 10 times

**arpit_dataguy** Highly Voted 4 years, 9 months ago

Selected Answer: BC

From snowflake documentation :

You might want to consider disabling auto-suspend for a warehouse if:

You have a heavy, steady workload for the warehouse.

You require the warehouse to be available with no delay or lag time. Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer.

upvoted 10 times

**LautaroBarone** 3 years, 6 months ago

Source: https://docs.snowflake.com/en/user-guide/warehouses-considerations#automating-warehouse-suspension

upvoted 2 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B and C are correct

upvoted 1 times

**Amitsnowflake** 3 years, 2 months ago

You might want to consider disabling auto-suspend for a warehouse if: You have a heavy, steady workload for the warehouse. You require the warehouse to be available with no delay or lag time.

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

BC correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

BC are correct

upvoted 1 times

Question #77 *Topic 1*

Which of the following are valid approaches to loading data into a Snowflake table? (Choose all that apply.)

A. Bulk copy from an External Stage

B. Continuous load using Snowpipe REST API

C. The Snowflake Web Interface (UI) data loading wizard

D. Bulk copy from an Internal Stage

**Suggested Answer:** *A, B, C*

*Community vote distribution*

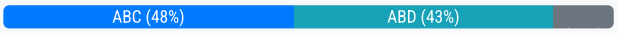


**redsky0** Highly Voted 4 years, 10 months ago

All are correct.

upvoted 36 times

**baig123** Highly Voted 3 years, 10 months ago

ABCD are correct, data can be loaded into snowflake table from internal stage using bulk copying.All are valid but the question said "valid approaches" so as approach and not as spot operation, i think that the "The Snowflake Web Interface (UI) data loading wizard" is not a valid approach. Do you agree with me?well its hard to say its not valid but logically better not to use it this way ?

upvoted 13 times

**baig123** 3 years, 10 months ago

ABCD are correct, data can be loaded into snowflake table from internal stage using bulk copying.

upvoted 13 times

**BigDataBB** 3 years, 10 months ago

All are valid but the question said "valid approaches" so as approach and not as spot operation, i think that the "The Snowflake Web Interface (UI) data loading wizard" is not a valid approach. Do you agree with me?

upvoted 3 times

**BigDataBB** 3 years, 10 months ago

All are valid but the question said "valid approaches" so as approach and not as spot operation, i think that the "The Snowflake Web Interface (UI) data loading wizard" is not a valid approach. Do you agree with me?well its hard to say its not valid but logically better not to use it this way ?

upvoted 3 times

**Hopefan** 3 years, 5 months ago

well its hard to say its not valid but logically better not to use it this way ?

upvoted 1 times

**niho1993** Most Recent 1 year, 1 month ago

A. Bulk copy from an External Stage  You can stage files in Amazon S3, Azure Blob, or GCS and load with COPY INTO. B. Continuous load using Snowpipe REST API  Snowpipe provides serverless, continuous loading of staged data files (triggered by REST API or cloud messaging). C. The Snowflake Web Interface (UI) data loading wizard  The UI wizard lets you load small files interactively — mainly for quick testing or ad-hoc loads. D. Bulk copy from an Internal Stage  You can stage files in user, table, or named internal stages and load into Snowflake using COPY INTO.

upvoted 1 times

**dfd8f9f** 1 year, 6 months ago

all are correct.

upvoted 1 times

**jeny12138** 1 year, 7 months ago

all correct

upvoted 1 times

**MultiCloudIronMan** 1 year, 10 months ago

All are correct

upvoted 2 times

**Shipra123** 1 year, 11 months ago

All are correct options to load data.

upvoted 2 times

**iqraaayyy** 2 years ago

All are correct https://docs.snowflake.com/en/user-guide/data-load-web-ui

upvoted 2 times

**DataEngineerNeeraj1** 2 years, 1 month ago

All are correct

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

All options

upvoted 2 times

**pranalig** 2 years, 4 months ago

correct answer:ABCD

upvoted 2 times

**PrashantGupta1616** 2 years, 5 months ago

The valid approaches to loading data into a Snowflake table are: A. Bulk copy from an External Stage B. Continuous load using Snowpipe REST API D. Bulk copy from an Internal Stage Therefore, the correct options are: A. Bulk copy from an External Stage B. Continuous load using Snowpipe REST API D. Bulk copy from an Internal Stage

upvoted 2 times

**BobFar** 3 years ago

ABCD is correct!

upvoted 2 times

**Roohee** 3 years, 1 month ago

bulk loading from internal stage is also possible. A B C D all are correct and supprted by snowflake https://docs.snowflake.com/en/user-guide/data-load-local-file-system

upvoted 3 times

**lofibeats** 3 years, 2 months ago

correct

upvoted 2 times

**MultiCloudIronMan** 3 years, 3 months ago

ABCD Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

ABC are correct

upvoted 1 times

Question #78 *Topic 1*

If auto-suspend is enabled for a Virtual Warehouse, the Warehouse is automatically suspended when:

A. All Snowflakes sessions using the Warehouse are terminated.

B. The last query using the Warehouse completes.

C. There are no users logged into Snowflake.

D. The Warehouse is inactive for a specified period of time.

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**Wzghaer** Most Recent 10 months, 1 week ago

By default, auto-suspend is enabled. Snowflake automatically suspends the warehouse if it is inactive for the specified period of time. By default, auto-resume is enabled. Snowflake automatically resumes the warehouse when any statement that requires a warehouse is submitted and the warehouse is the current warehouse for the session.

upvoted 1 times

**hotflake** 2 years, 3 months ago

Auto suspend turns off a virtual warehouse after a specified amount of idle time.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

D is correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

D is correct

upvoted 1 times

**SV1122** 3 years, 10 months ago

Answer is D https://docs.snowflake.com/en/user-guide/warehouses-overview.html#auto-suspension-and-auto-resumption

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

D is correct

upvoted 2 times

**kaun** 4 years, 1 month ago

Hi could you please add or update it some more questions. I took Snowflake pro core exam last week but didn't able to pass the exam. I have gone all the questions which is here but even one question didn't come to the exam from this.

upvoted 3 times

**kjdfsafc87456** 4 years, 2 months ago

The answer is correct

upvoted 1 times

**MinuteMaid** 4 years, 7 months ago

The correct answer is D

upvoted 3 times

Question #79 *Topic 1*

True or False: Multi-Factor Authentication (MFA) in Snowflake is only supported in conjunction with Single Sign-On (SSO).

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B is Correct

upvoted 1 times

**SushY18** 3 years, 5 months ago

MFA token caching can be combined with connection caching in federated single sign-on. To combine these two features, ensure that the ALLOW_ID_TOKEN parameter is set to true in tandem with the ALLOW_CLIENT_MFA_CACHING parameter

upvoted 3 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**EmiB** 3 years, 7 months ago

B. False

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

B is correct

upvoted 1 times

**moon2351** 4 years, 10 months ago

Answer is correct.

upvoted 3 times

Question #80 *Topic 1*

The number of queries that a Virtual Warehouse can concurrently process is determined by (Choose two.):

A. The complexity of each query

B. The CONCURRENT_QUERY_LIMIT parameter set on the Snowflake account

C. The size of the data required for each query

D. The tool that is executing the query

**Suggested Answer:** *A, C*

*Community vote distribution*

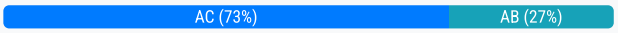


**DataEngineerNeeraj1** Highly Voted 2 years, 1 month ago

Only A and C is correct , B would also have been corrected if MAX_CONCURRENCY_LEVEL was given but here option is MAX_CONCURRENCY_LIMIT .

upvoted 6 times

**dfd8f9f** Most Recent 1 year, 6 months ago

a & b are the answers

upvoted 1 times

**p22_nilesh** 2 years, 5 months ago

AB is the answer

upvoted 2 times

**ifmateus** 3 years, 2 months ago

I think it's A and B because "You can use the MAX_CONCURRENCY_LEVEL parameter to limit the number of concurrent queries running in a warehouse." and "The number of queries that a warehouse can concurrently process is determined by the size and complexity of each query. " here it seems the number of queries processed are determined by the query complexity and its size and not the data size.https://docs.snowflake.com/en/user-guide/warehouses-overview#query-processing-and-concurrencyhttps://docs.snowflake.com/en/user-guide/performance-query-warehouse-max-concurrencyI'd agree, but the parameter in that case would be MAX_CONCURRENCY_LEVEL, not CONCURRENT_QUERY_LIMIT

upvoted 2 times

**ifmateus** 3 years, 2 months ago

I think it's A and B because "You can use the MAX_CONCURRENCY_LEVEL parameter to limit the number of concurrent queries running in a warehouse." and "The number of queries that a warehouse can concurrently process is determined by the size and complexity of each query. " here it seems the number of queries processed are determined by the query complexity and its size and not the data size.

upvoted 2 times

**ifmateus** 3 years, 2 months ago

https://docs.snowflake.com/en/user-guide/warehouses-overview#query-processing-and-concurrency

upvoted 2 times

**ifmateus** 3 years, 2 months ago

https://docs.snowflake.com/en/user-guide/warehouses-overview#query-processing-and-concurrencyhttps://docs.snowflake.com/en/user-guide/performance-query-warehouse-max-concurrencyI'd agree, but the parameter in that case would be MAX_CONCURRENCY_LEVEL, not CONCURRENT_QUERY_LIMIT

upvoted 3 times

**ifmateus** 3 years, 2 months ago

https://docs.snowflake.com/en/user-guide/performance-query-warehouse-max-concurrency

upvoted 1 times

**ifmateus** 3 years, 2 months ago

https://docs.snowflake.com/en/user-guide/performance-query-warehouse-max-concurrencyI'd agree, but the parameter in that case would be MAX_CONCURRENCY_LEVEL, not CONCURRENT_QUERY_LIMIT

upvoted 1 times

**oscarglob** 3 years ago

I'd agree, but the parameter in that case would be MAX_CONCURRENCY_LEVEL, not CONCURRENT_QUERY_LIMIT

upvoted 4 times

**Mayuri917233** 3 years, 5 months ago

AC is correct

upvoted 1 times

**systematics78** 3 years, 9 months ago

https://docs.snowflake.com/en/user-guide/warehouses-overview.html#query-processing-and-concurrency AC are correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

AC are correct

upvoted 3 times

**vivek_nitd** 4 years, 8 months ago

Answer: AC MAX_CONCURRENCY_LEVEL is set to 8 by default at account level and can be changed. In reality, we can have more than 8 concurrent queries as well. But it depends on the factors such as the complexity of the queries, their resource consumptions etc.

upvoted 4 times

**PrithviPatil** 4 years, 10 months ago

What about parameter MAX_CONCURRENCY_LEVEL that defines the maximum number of parallel or concurrent statements a warehouse can execute? https://community.snowflake.com/s/article/Warehouse-Concurrency-and-Statement-Timeout-Parameters

upvoted 1 times

**moon2351** 4 years, 10 months ago

Answer is correct. The number of queries that a warehouse can concurrently process is determined by the size and complexity of each query. As queries are submitted, the warehouse calculates and reserves the compute resources needed to process each query. If the warehouse does not have enough remaining resources to process a query, the query is queued, pending resources that become available as other running queries complete.

upvoted 4 times

Question #81 *Topic 1*

Which of the following statements are true of VALIDATION_MODE in Snowflake? (Choose two.)

A. The VALIDATION_MODE option is used when creating an Internal Stage

B. VALIDATION_MODE=RETURN_ALL_ERRORS is a parameter of the COPY command

C. The VALIDATION_MODE option will validate data to be loaded by the COPY statement while completing the load and will return the rows that could not be loaded without error

D. The VALIDATION_MODE option will validate data to be loaded by the COPY statement without completing the load and will return possible errors

**Suggested Answer:** *B, D*

*Community vote distribution*

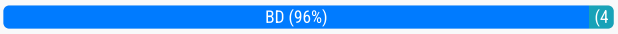


**pranalig** Most Recent 2 years, 4 months ago
correct answer:BD
upvoted 2 times

**trainingRalph** 3 years ago
The VALIDATION_MODE copy option instructs a COPY statement to validate the data to be loaded and return results based on the validation option specified. No data is loaded when this copy option is specified. For more information about the copy option, see COPY INTO .
https://docs.snowflake.com/en/user-guide/data-load-bulk-ts#step-2-validating-the-data-load
upvoted 2 times

**Thaiya** 3 years, 1 month ago
B & D
upvoted 3 times

**SamyakSahoo** 3 years, 4 months ago
BD are correct
upvoted 2 times

**malexandersalazar** 3 years, 5 months ago
BD, the COPY command only tests the files for errors.
upvoted 1 times

**Mayuri917233** 3 years, 5 months ago
BC are correct
upvoted 1 times

**isolated_sniper** 3 years, 6 months ago
B and D are correct. Validating the Data Load The VALIDATION_MODE copy option instructs a COPY statement to validate the data to be loaded and return results based on the validation option specified. No data is loaded when this copy option is specified. For more information about the copy option, see COPY INTO . ref: https://docs.snowflake.com/en/user-guide/data-load-bulk-ts
upvoted 1 times

**RajuNaik** 3 years, 6 months ago
BD are correct
upvoted 1 times

**Ayush24** 3 years, 6 months ago
VALIDATION_MODE = RETURN_n_ROWS | RETURN_ERRORS | RETURN_ALL_ERRORS String (constant) that instructs the COPY command to validate the data files instead of loading them into the specified table; i.e. the COPY command tests the files for errors but does not load them.
upvoted 3 times

**nencheppan** 3 years, 7 months ago
BD are correct
upvoted 1 times

**Priye** 3 years, 8 months ago

BD are correct

upvoted 1 times

**ajay_1233456** 3 years, 9 months ago

BD are correct

upvoted 1 times

**SV1122** 3 years, 10 months ago

Was on exam Dec 31st, 2022

upvoted 3 times

**baig123** 3 years, 10 months ago

B & D are correct answer.

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

BD are correct

upvoted 2 times

**bharathmv** 4 years, 1 month ago

B and D is right

upvoted 2 times

**kjdfsafc87456** 4 years, 2 months ago

there won't be any data loaded The VALIDATION_MODE copy option instructs a COPY statement to validate the data to be loaded and return results based on the validation option specified. No data is loaded when this copy option is specified.

upvoted 4 times

Question #82 *Topic 1*

What privileges are required to create a task?

A. The GLOBAL privilege CREATE TASK is required to create a new task.

B. Tasks are created at the Application level and can only be created by the Account Admin role.

C. Many Snowflake DDLs are metadata operations only, and CREATE TASK DDL can be executed without virtual warehouse requirement or task specific grants.

D. The role must have access to the target schema and the CREATE TASK privilege on the schema itself.

**Suggested Answer:** *D*

*Community vote distribution*




**shubtred** Highly Voted 3 years, 4 months ago

D. The role must have access to the target schema and the CREATE TASK privilege on the schema itself: The role executing the CREATE TASK statement must have access privileges on the schema where the task will be created. This includes having the necessary privileges to create objects within the schema. Typically, the role needs the CREATE TASK privilege granted on the specific schema.

upvoted 7 times

**pranalig** Most Recent 2 years, 4 months ago

correct answer:D

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

D. The role must have access to the target schema and the CREATE TASK privilege on the schema itself. In Snowflake, to create a task, the role must have access to the target schema where the task will be created, and it needs the CREATE TASK privilege on that schema. This ensures that users can only create tasks in schemas where they have been granted the necessary permissions

upvoted 1 times

**krishnak2244** 3 years, 5 months ago

Ans is D Global Privileges dont have CREATE TASK in it CREATE TASK is present in Schema level privileges only Doc: https://docs.snowflake.com/en/user-guide/security-access-control-privileges

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**SV1122** 3 years, 10 months ago

Was on exam Dec 31st, 2022so many comments with Dec 31st, 2022 , please change date.

upvoted 3 times

**SV1122** 3 years, 10 months ago

Was on exam Dec 31st, 2022

upvoted 3 times

**jw1806** 3 years, 7 months ago

so many comments with Dec 31st, 2022 , please change date.

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

D is correct

upvoted 1 times

**bharathmv** 4 years, 1 month ago

D is answer

upvoted 1 times

**arunshai** 4 years, 9 months ago

Ans - D

upvoted 1 times

**Sid460545** 4 years, 9 months ago

D is correct

upvoted 1 times

**arpit_dataguy** 4 years, 9 months ago

Only CREATE TASK at the schema level is required, not global. All tasks in a simple tree must have the same task owner (i.e. a single role must have the OWNERSHIP privilege on all of the tasks in the tree). All tasks in a simple tree must exist in the same schema.

upvoted 4 times

**redsky0** 4 years, 9 months ago

Only CREATE TASK at the schema level is required, not global: https://docs.snowflake.com/en/user-guide/tasks-intro.html#creating-tasks

upvoted 4 times

Question #83 *Topic 1*

What are the three things customer want most from their enterprise data warehouse solution? (Choose three.)

A. On-premise availability

B. Simplicity

C. Open source based

D. Concurrency

E. Performance

**Suggested Answer:** *B, D, E*

*Community vote distribution*

BDE (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B,D,E are Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

BDE are correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

BDE is correct

upvoted 3 times

**Antz83** 4 years, 3 months ago

Snowflake Marketing Slide from their training. Definitely Simplicity, Concurrency and Performance is consistently mentioned.

upvoted 2 times

Question #84 *Topic 1*

True or False: Some queries can be answered through the metadata cache and do not require an active Virtual Warehouse.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*

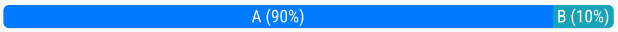


**blazingliverbird** Highly Voted 4 years, 10 months ago

Some aggregate queries are answered thru micro partitions metadata only not requiring any VW spin ups.

upvoted 7 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Yes,It is True

upvoted 1 times

**pranalig** 2 years, 4 months ago

correct answer:A

upvoted 1 times

**Gavin1272** 2 years, 8 months ago

Metadata queries can be done as metadata is stored in the services layer and does not require the compute layer to spin up a vm to access the storage layer.

upvoted 1 times

**MahDeep** 3 years, 1 month ago

Since Metadata cache and the Query Result cache are part of the Cloud Services Layer, they can run without the need of an active virtual warehouse

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**isolated_sniper** 3 years, 6 months ago

Answer is A.

upvoted 4 times

**avijit1982** 3 years, 6 months ago

Sorry, just tested it. It is A.

upvoted 3 times

**avijit1982** 3 years, 6 months ago

Actually, I think one needs active datawarehouse to run the query but it would not use it.

upvoted 1 times

**alanhuangjob** 3 years, 7 months ago

There is metadata cache in the service layer, so answer A is correct.

upvoted 1 times

**KarBiswa** 3 years, 9 months ago

Count min does not require

upvoted 2 times

**12hj** 3 years, 9 months ago

Answer - B, because metadata has no cache.

upvoted 1 times

**Chandra720** 3 years, 10 months ago

If there is no change in query then it will fetch from metadata cache I.e from cloud service layer.

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

A is correct

upvoted 2 times

**jveag** 4 years, 1 month ago

should be A

upvoted 2 times

**viratang** 4 years, 4 months ago

Below query doesn't need an active warehouse to run for a seond time. SELECT COUNT(*) FROM "SNOWFLAKE_SAMPLE_DATA"."TPCH_SF1"."CUSTOMER";

upvoted 2 times

**BungyTex** 4 years, 4 months ago

its A as metadata has aggregate data already

upvoted 2 times

Question #85 *Topic 1*

When scaling out by adding clusters to a multi-cluster warehouse, you are primarily scaling for improved:

A. Concurrency

B. Performance

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**marklv** Most Recent 2 years, 1 month ago

Improving concurrency would also improve performance. Primarily is the key word here.

upvoted 1 times

**Karthikparasuraman** 2 years, 2 months ago

https://docs.snowflake.com/en/user-guide/warehouses-considerations#scaling-up-vs-scaling-out

upvoted 1 times

**MultiCloudIronMan** 2 years, 3 months ago

Correct

upvoted 1 times

**Mayuri917233** 2 years, 5 months ago

A is correct

upvoted 1 times

**SteveCrown** 2 years, 6 months ago

Answer: A A. Concurrency is the primary benefit when scaling out by adding clusters to a multi-cluster warehouse. By adding more clusters, you increase the number of compute resources available, which allows the warehouse to handle more queries and users concurrently. This can help prevent bottlenecks and improve overall query performance, but the primary goal of adding clusters is to increase concurrency. B. Performance can also improve as a result of scaling out, but it is not the primary benefit. The main focus is on improving concurrency by enabling more users and queries to be processed simultaneously.

upvoted 3 times

**fahfouhi94** 2 years, 8 months ago

A is correct

upvoted 1 times

**SV1122** 2 years, 10 months ago

Concurrency is correct

upvoted 1 times

**lhstudylh** 3 years, 8 months ago

Correct answer

upvoted 2 times

Question #86 *Topic 1*

What is the minimum Snowflake edition that provides data sharing?

A. Standard

B. Premier

C. Enterprise

D. Business Critical Edition

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**someshVashistha** Most Recent 1 year, 10 months ago

A.Standard

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A.Standard

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

Data sharing is possible in the Standard Edition also. VPS (Virtual Private Snowflake) does not support Secure Data Sharing due to the current limitations against sharing data across regions. Standard and Enterprise Editions support Secure Data Sharing with the usual caveats.

upvoted 2 times

**lhstudylh** 4 years, 8 months ago

Correct answer

upvoted 2 times

Question #87 *Topic 1*

True or False: Each worksheet in the Snowflake Web Interface (UI) can be associated with different roles, databases, schemas, and Virtual Warehouses.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*

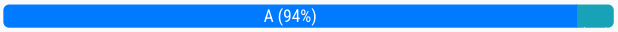


**Prajdeo** Highly Voted 3 years, 8 months ago

This is True. Correct answer should be A https://docs.snowflake.com/en/user-guide/ui-worksheet Note Your current interface role determines the default role for worksheets that you open, but the worksheets are not tied to the interface role. Each worksheet has its own role that can be set independently.

upvoted 6 times

**Mallikharjuna452** 2 years, 3 months ago

Yes,It is True.

upvoted 1 times

**pranalig** 2 years, 4 months ago

correct answer:A

upvoted 1 times

**Alex_Ovidiu** 2 years, 7 months ago

False: Each worksheet in the Snowflake Web Interface (UI) cannot be associated with different roles, databases, schemas, and Virtual Warehouses. Here's why: The Snowflake Web UI maintains a single connection context at a time. This means you can only connect using a specific role, database, schema, and virtual warehouse for each session. While the UI allows you to switch between different roles, databases, schemas, and virtual warehouses, these changes apply globally to all your worksheets within that session. Some data manipulation tools might offer functionalities where each worksheet can have independent connections or contexts. However, the Snowflake Web UI itself doesn't support such a feature.This statement from your thread above is simply incorrect. "While the UI allows you to switch between different roles, databases, schemas, and virtual warehouses, these changes apply globally to all your worksheets within that session." You can change all of these parameters at any time for any worksheet.

upvoted 1 times

**Alex_Ovidiu** 2 years, 7 months ago

False: Each worksheet in the Snowflake Web Interface (UI) cannot be associated with different roles, databases, schemas, and Virtual Warehouses.

Here's why:

The Snowflake Web UI maintains a single connection context at a time. This means you can only connect using a specific role, database, schema, and virtual warehouse for each session. While the UI allows you to switch between different roles, databases, schemas, and virtual warehouses, these changes apply globally to all your worksheets within that session. Some data manipulation tools might offer functionalities where each worksheet can have independent connections or contexts. However, the Snowflake Web UI itself doesn't support such a feature.

upvoted 1 times

**56d43d6** Most Recent 1 year, 10 months ago

This statement from your thread above is simply incorrect. "While the UI allows you to switch between different roles, databases, schemas, and virtual warehouses, these changes apply globally to all your worksheets within that session." You can change all of these parameters at any time for any worksheet.

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

Selected Answer: B

B is correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correctIncorrect, a worksheet uses the following USE warehouse; USE role USE schema USE DB

upvoted 1 times

**bbrandhh** Highly Voted 3 years, 3 months ago

Incorrect, a worksheet uses the following USE warehouse; USE role USE schema USE DB

upvoted 5 times

**ajay_1233456** 3 years, 9 months ago

A is correct as each worksheet can be associated with different roles

upvoted 1 times

**baig123** 3 years, 10 months ago

A is correct answer. we can set different context parameters at different worksheet.

upvoted 2 times

**Tapasgup007** 3 years, 11 months ago

A is correct as each worksheet can be associated with different roles

upvoted 1 times

**jveag** 4 years, 1 month ago

these tests are very confusing.these tests are very confusing. I guess it is false because the question means that the worksheets are not bound to a certain role, user...etc.. can change without losing the work done. https://docs.snowflake.com/en/user-guide/ui-worksheet.html Dropdown menu: Change the current database, schema, or warehouse for the current worksheet without losing your work. Resume/suspend or resize your current warehouse.

upvoted 2 times

**jveag** 4 years, 1 month ago

Selected Answer: A

these tests are very confusing.

upvoted 2 times

**jveag** 4 years, 1 month ago

these tests are very confusing. I guess it is false because the question means that the worksheets are not bound to a certain role, user...etc.. can change without losing the work done. https://docs.snowflake.com/en/user-guide/ui-worksheet.html Dropdown menu: Change the current database, schema, or warehouse for the current worksheet without losing your work. Resume/suspend or resize your current warehouse.

upvoted 2 times

**kjdfsafc87456** 4 years, 2 months ago

It would be true!

upvoted 1 times

**distopic** 4 years, 4 months ago

A is the answer

upvoted 2 times

**learnazureportal** 4 years, 6 months ago

Answer is A

upvoted 2 times

**CNBOOST2** 4 years, 7 months ago

Answer is B I think

upvoted 1 times

**Rajokkiyam** 4 years, 8 months ago

Answer Truewich answeR?

upvoted 2 times

**Rajokkiyam** 4 years, 8 months ago

Answer True

upvoted 2 times

**CNBOOST2** 4 years, 7 months ago

wich answeR?

upvoted 1 times

**Sid460545** 4 years, 9 months ago

Answer is True

upvoted 1 times

**Casual** 4 years, 9 months ago

Answer is A

upvoted 1 times

**blazingliverbird** 4 years, 10 months ago

Each worksheet can be associated with different parameters.

upvoted 3 times

Question #88 *Topic 1*

True or False: You can query the files in an External Stage directly without having to load the data into a table.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**PrasantSadangi** Most Recent 1 year, 7 months ago

Querying Data in Staged Files Snowflake supports using standard SQL to query data files located in an internal (i.e. Snowflake) stage or named external (Amazon S3, Google Cloud Storage, or Microsoft Azure) stage. This can be useful for inspecting/viewing the contents of the staged files, particularly before loading or after unloading data. https://docs.snowflake.com/en/user-guide/querying-stage

upvoted 2 times

**Alex_Ovidiu** 2 years, 1 month ago

False: You cannot directly query the files in an External Stage without having to load the data into a table in Snowflake. Here's why: External stages act as pointers to data residing in external cloud storage locations like S3 or Azure Blob Storage. Snowflake doesn't have direct access to the data format or schema within those files. To query the data, you need to first load it into a Snowflake table using the COPY INTO command. This process reads the data from the external stage, parses it based on the specified format definition, and populates the target table within Snowflake. Once the data is loaded into a Snowflake table, you can then use standard SQL queries to access and analyze it. In essence, external stages serve as a staging area for data before it's brought into Snowflake for processing and querying.

upvoted 3 times

**Bujji1234** 2 years, 3 months ago

A is correct

upvoted 1 times

**oscarglob** 2 years, 5 months ago

True: https://docs.snowflake.com/en/user-guide/querying-stage

upvoted 2 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**Mayuri917233** 2 years, 11 months ago

A is correct

upvoted 1 times

**Tapasgup007** 3 years, 5 months ago

A is correct

upvoted 1 times

**kjdfsafc87456** 3 years, 8 months ago

Snowflake supports using standard SQL to query data files located in an internal (i.e. Snowflake) stage or named external (Amazon S3, Google Cloud Storage, or Microsoft Azure) stage. This can be useful for inspecting/viewing the contents of the staged files, particularly before loading or after unloading data.Yes, source: https://docs.snowflake.com/en/user-guide/querying-stage

upvoted 2 times

**kjdfsafc87456** 3 years, 8 months ago

Snowflake supports using standard SQL to query data files located in an internal (i.e. Snowflake) stage or named external (Amazon S3, Google Cloud Storage, or Microsoft Azure) stage. This can be useful for inspecting/viewing the contents of the staged files, particularly before loading or after unloading data.

upvoted 2 times

**oscarglob** 2 years, 5 months ago

Yes, source: https://docs.snowflake.com/en/user-guide/querying-stage

upvoted 1 times

**lhstudylh** 4 years, 2 months ago

Correct answer External tables are read-only, therefore no DML operations can be performed on them; however, external tables can be used for query and join operations. Views can be created against external tables.Yeah, but that wasn't the question... this one is about whether you can query files on an external stage directly

upvoted 1 times

**lhstudylh** 4 years, 2 months ago

**Selected Answer: A**

Correct answer

External tables are read-only, therefore no DML operations can be performed on them; however, external tables can be used for query and join operations. Views can be created against external tables.

upvoted 1 times

**oscarglob** 2 years, 5 months ago

Yeah, but that wasn't the question... this one is about whether you can query files on an external stage directly

upvoted 1 times

Question #89 *Topic 1*

The FLATTEN function is used to query which type of data in Snowflake?

A. Structured data

B. Semi-structured data

C. Both of the above

D. None of the above

**Suggested Answer:** *C*

*Community vote distribution*

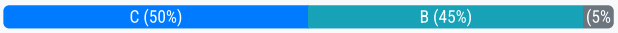


**alexhparreira** Most Recent 5 months ago

Why Option B is correct: The FLATTEN function is a table function specifically designed to unpack and explode nested, semi-structured data types (such as VARIANT, OBJECT, or ARRAY) into separate lateral rows. It converts hierarchical data (like JSON or XML) into a relational format.

upvoted 1 times

**francisss** 1 year, 8 months ago

* Why is FLATTEN also used in Structured Types? : Because Snowflake's Structured Types (structured data) have internally nested complex structures, they are difficult to query directly using general SQL queries. : At this time, the FLATTEN function performs the following roles: - Convert nested structures into table form for easy search. - Internal composite elements (STRUCTURED ARRAY, STRUCTURED OBJECT, MAP internal elements) are expanded into separate rows to enable querying. In other words, the FLATTEN function is essential in Snowflake's Structured Types (explicitly structured complex data), not general table data.

upvoted 1 times

**francisss** 1 year, 8 months ago

there are 2 reason in snowflake document. 1) https://docs.snowflake.com/en/sql-reference/data-types-structured#label-structured-types-working-flatten --> Using the FLATTEN function with values of structured types --> You can pass structured ARRAY, structured OBJECT, and MAP values to the FLATTEN function. As is the case with semi-structured data types, you can use the PATH argument to specify the value being flattened. 2) https://docs.snowflake.com/en/sql-reference/functions/flatten#usage-notes --> FLATTEN is a table function that takes a VARIANT, OBJECT, or ARRAY column and produces a lateral view (i.e. an inline view that contains correlation referring to other tables that precede it in the FROM clause). --> FLATTEN can be used to convert semi-structured data to a relational representation.

upvoted 1 times

**someshVashistha** 1 year, 10 months ago

Answer B The FLATTEN function in Snowflake is used to query and extract data from semi-structured data types, specifically: Arrays Objects (key-value pairs, such as those found in JSON data) The function expands these structures into a table-like format, allowing you to work with individual elements or key-value pairs in SQL queries.

upvoted 1 times

**MohamedBen** 1 year, 11 months ago

Answer B

upvoted 1 times

**DataEngineerNeeraj1** 2 years, 1 month ago

C

upvoted 1 times

**aemilka** 2 years, 2 months ago

You can pass structured ARRAYs, structured OBJECTs, and MAPs to the FLATTEN function. https://docs.snowflake.com/en/sql-reference/data-types-structured#label-structured-types-working-flatten FLATTEN can be used to convert semi-structured data to a relational representation. https://docs.snowflake.com/en/sql-reference/functions/flatten So the answer is C. Both of the above

upvoted 2 times

**omiiUST** 2 years, 3 months ago

The main purpose of FLATTEN is to Un-nest the semistructured data like JSON.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

It works for both

upvoted 1 times

**372110c** 2 years, 4 months ago

Works with both

upvoted 2 times

**MBrat** 2 years, 5 months ago

C is correct, see https://docs.snowflake.com/en/sql-reference/data-types-structured#label-structured-types-working-flatten

upvoted 3 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

B is correct

upvoted 1 times

**lhstudylh** 4 years, 8 months ago

Correct answer FLATTEN is used to unnest semi-structured data. Don't see an application for structured data as by definition it shouldn't be nested.

upvoted 3 times

Question #90 *Topic 1*

True or False: An active warehouse is required to run a COPY INTO statement.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*




**nkit** Most Recent 1 year, 7 months ago

No, an active warehouse is not required to run a COPY INTO statement in Snowflake. The COPY INTO command can be executed without an active warehouse when: Loading data from external stages (like Amazon S3, Google Cloud Storage, or Microsoft Azure) Using the AUTO_COMPRESS option

upvoted 2 times

**hotflake** 2 years, 3 months ago

COPY INTO: ● Copies data from staged files to a table ● All files loaded in a single operation must use the same stage and file format ● Can auto-detect compression mode, date and time formats ● Requires an active virtual warehouse

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A.True,while loading the data,an active warehouse is required

upvoted 1 times

**Alex_Ovidiu** 2 years, 7 months ago

False: An active warehouse is not always required to run a COPY INTO statement in Snowflake. Here's why: The COPY INTO statement is used to load data from external locations (stages) into Snowflake tables. While virtual warehouses provide the compute resources for processing queries within Snowflake, the COPY INTO operation itself is a data movement process.

upvoted 3 times

**Rajivnb** 3 years ago

COPY INTO table basically rights data into Micro Partition. To access the micro partition it requires a active warehouse to pass thru right. To access the storage it should pass thru the warehouse. it should be active. even if its suspended, upon receiving a Query the warehouse will resume and come into active to execute the query.

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**Raju039** 3 years, 8 months ago

https://docs.snowflake.com/en/user-guide/getting-started-tutorial-copy-into

upvoted 2 times

**markpumc** 3 years, 9 months ago

Should be false B: False. A COPY INTO statement in Snowflake does not require an active warehouse to run. The statement is used to load data into a Snowflake table from a file stored in a cloud-based object storage such as Amazon S3, Microsoft Azure Blob Storage, or Google Cloud Storage. An active warehouse is only required to perform computationally-intensive tasks such as data processing, querying, and data transformation.

upvoted 1 times

**markpumc** 3 years, 9 months ago

Should be false B: False. A COPY INTO statement in Snowflake does not require an active warehouse to run. The statement is used to load data into a Snowflake table from a file stored in a cloud-based object storage such as Amazon S3, Microsoft Azure Blob Storage, or Google Cloud Storage. An active warehouse is only required to perform computationally-intensive tasks such as data processing, querying, and data transformation.COPY INTO table basically rights data into Micro Partition. To access the micro partition it requires a active warehouse to pass thru right. To access the storage it should pass thru the warehouse. it should be active. even if its suspended, upon receiving a Query the warehouse will resume and come into active to execute the query.

upvoted 1 times

**Tapasgup007** 3 years, 11 months ago

A is correct

upvoted 1 times

**kjdfsafc87456** 4 years, 2 months ago

Execute COPY INTO to load your staged data into the target table. Note that this command requires an active, running warehouse, which you created as a prerequisite for this tutorial. If you don't have access to a warehouse, you will need to create one now.

upvoted 3 times

**Sly007** 4 years, 6 months ago

Answer is A. Execute COPY INTO to load your staged data into the target table. Note that this command requires an active, running warehouse. https://docs.snowflake.com/en/user-guide/getting-started-tutorial-copy-into.html

upvoted 2 times

**learnazureportal** 4 years, 6 months ago

Answer is B

upvoted 1 times

**learnazureportal** 4 years, 6 months ago

Answer is BCorrected = A is Answer

upvoted 1 times

**learnazureportal** 4 years, 6 months ago

Corrected = A is Answer

upvoted 1 times

**addixion** 4 years, 7 months ago

I think it is false

upvoted 1 times

**lhstudylh** 4 years, 8 months ago

Correct answer Link below states that you need an active, running warehouse to before COPY INTO functions https://docs.snowflake.com/en/user-guide/getting-started-tutorial-copy-into.html

upvoted 1 times

Question #91 *Topic 1*

True or False: AWS Private Link provides a secure connection from the Customer's on-premise data center to the Snowflake.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

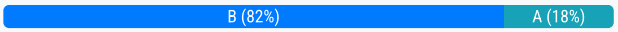


**pappuji** Highly Voted 4 years, 4 months ago

n addition, if you have an on-premises environment (e.g. a non-hosted data center), you can choose to use AWS Direct Connect, in conjunction with AWS PrivateLink, to connect all your virtual and physical environments in a single, private network. therefore AWS private link alone wont work.
Answer:False

upvoted 6 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B.False,Aws Private link provides a secure connection between Vpc's and AWS Services

upvoted 1 times

**titihab** 3 years ago

From SnowFlake Documentation : AWS PrivateLink & Snowflake BUSINESS CRITICAL FEATURE

This feature requires Business Critical (or higher).

This topic describes how to configure AWS PrivateLink to directly connect your Snowflake account to one or more AWS VPCs.

Note that AWS PrivateLink is not a service provided by Snowflake. It is an AWS service that Snowflake supports to use with your Snowflake account.

upvoted 1 times

**titihab** 3 years ago

From SnowFlake Documentation : AWS PrivateLink & Snowflake BUSINESS CRITICAL FEATURE This feature requires Business Critical (or higher). This topic describes how to configure AWS PrivateLink to directly connect your Snowflake account to one or more AWS VPCs. Note that AWS PrivateLink is not a service provided by Snowflake. It is an AWS service that Snowflake supports to use with your Snowflake account.then it's B

upvoted 1 times

**titihab** 3 years ago

then it's B

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**alanhuangjob** 3 years, 7 months ago

AWS Direct connect for on-premise to Snowflake VPC; AWS Private Link for different AWS VPCs in the same AWS region.

upvoted 4 times

**Anch_D** 3 years, 7 months ago

AWS VPC provide secure direct connection between clients AWS VPC and Snowflake VPC. not the onprim

upvoted 4 times

**Tapasgup007** 3 years, 10 months ago

B is correct

upvoted 1 times

**distopic** 4 years, 4 months ago

AWS private link is not enough, i think that the answer is b

upvoted 4 times

**DK69** 4 years, 5 months ago

Answer is B

upvoted 1 times

**learnazureportal** 4 years, 6 months ago

Answer is A

upvoted 1 times

**Casual** 4 years, 8 months ago

Answer is B if you have an on-premises environment, you can choose to use AWS Direct Connect

upvoted 4 times

**grabc** 4 years, 9 months ago

So it should be A?No , the correct Answer is B.

upvoted 1 times

**grabc** 4 years, 9 months ago

So it should be A?

upvoted 1 times

**BobFar** 3 years ago

No , the correct Answer is B.

upvoted 1 times

**arpit_dataguy** 4 years, 9 months ago

AWS PrivateLinkis an AWS service for creating private VPC endpoints that allow direct, secure connectivity between AWS VPCs without traversing the public Internet. Because Snowflake on AWS is implemented as a VPC, PrivateLink enables creating a highly-secure network between Snowflake and your other AWS VPCs (in the same AWS region), fully protected from unauthorized external access.No @arpit_dataguy is B because "AWS PrivateLinkis"-> private VPC between your AWS VPCs and the Snowflake VPC. In addition, if you have an on-premises environment (e.g. a non-hosted data center), you can choose to use AWS Direct Connect, in conjunction with AWS PrivateLink, to connect all your virtual and physical environments in a single, private network. So to be correct A is needed also "AWS Direct Connect"

upvoted 3 times

**arpit_dataguy** 4 years, 9 months ago

**Selected Answer: A**

AWS PrivateLinkis an AWS service for creating private VPC endpoints that allow direct, secure connectivity between AWS VPCs without traversing the public Internet. Because Snowflake on AWS is implemented as a VPC, PrivateLink enables creating a highly-secure network between Snowflake and your other AWS VPCs (in the same AWS region), fully protected from unauthorized external access.

upvoted 3 times

**BigDataBB** 3 years, 11 months ago

No @arpit_dataguy is B because "AWS PrivateLinkis"-> private VPC between your AWS VPCs and the Snowflake VPC. In addition, if you have an on-premises environment (e.g. a non-hosted data center), you can choose to use AWS Direct Connect, in conjunction with AWS PrivateLink, to connect all your virtual and physical environments in a single, private network. So to be correct A is needed also "AWS Direct Connect"

upvoted 3 times

Question #92 *Topic 1*

True or False: Snowflake's Global Services Layer gathers and maintains statistics on all columns in all micro-partitions.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*

A (60%) B (40%)

**alfredofmt** Highly Voted 3 years, 4 months ago

Selected Answer: B

The Global Services Layer in Snowflake is responsible for several tasks including authentication, infrastructure management, metadata management, query parsing and optimization, and access control1. However, the documentation does not specifically mention that it gathers and maintains statistics on all columns in all micro-partitions.

upvoted 7 times

**alfredofmt** Highly Voted 3 years, 4 months ago

The Global Services Layer in Snowflake is responsible for several tasks including authentication, infrastructure management, metadata management, query parsing and optimization, and access control1. However, the documentation does not specifically mention that it gathers and maintains statistics on all columns in all micro-partitions.collecting metadata is called statistics. It basically collects how many distinct values, count(*), Min,Max values etc., this is statistics

upvoted 7 times

**Rajivnb** 3 years ago

collecting metadata is called statistics. It basically collects how many distinct values, count(*), Min,Max values etc., this is statistics

upvoted 2 times

**lobr** Most Recent 1 year, 5 months ago

Snowflake's Global Services Layer manages metadata, which includes statistics about the data stored in micro-partitions. However, it does not gather and maintain statistics on all columns in all micro-partitions. Snowflake automatically collects metadata and statistics to optimize query performance. It does this continuously in the background, analyzing data as it's loaded and as queries are executed. While it gathers essential statistics, it's unnecessary and inefficient to maintain detailed statistics for every single column in every micro-partition. Snowflake focuses on collecting statistics that are most useful for query optimization, such as: Value ranges within columns. The number of distinct values (approximate). Information about NULL values.

upvoted 1 times

**kentucky13caa** 1 year, 11 months ago

The answer is A. True. Let's explain why this is true: Snowflake's Global Services Layer: - Automatically gathers statistics on ALL columns - Maintains statistics for ALL micro-partitions - This is done without DBA intervention - Used for query optimization Key points about statistics collection: - Continuous and automatic process - Covers every column in every micro-partition - No manual statistics gathering needed - Used by query optimizer for efficient execution plans For the SnowPro Core exam, remember: - Statistics collection is automatic - No DBA maintenance required - Covers all columns and micro-partitions - This is a key differentiator from traditional databases - Helps with query optimization and performance - Part of Snowflake's cloud-native architecture benefits This automatic statistics gathering is one of Snowflake's key features that helps with query optimization and maintenance-free operations.

upvoted 3 times

**hotflake** 2 years, 3 months ago

● Micro-partition level metadata ○ Row count ● Micro-partition column-level metadata ○ MIN and MAX values ○ Number of DISTINCT values ○ Number of NULL values

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B.False,Snowflake's Global Services Layer gathers and maintains statistics on columns in micro-partitions, but it does not gather and maintain statistics on all columns in all micro-partitions

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

Snowflake's Global Services Layer does not gather and maintain statistics on all columns in all micro-partitions. Instead, it dynamically gathers statistics as needed for query optimization, utilizing metadata and sampling techniques to optimize query performance.

upvoted 1 times

**Stevec80** 2 years, 5 months ago

https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions#what-are-micro-partitions

upvoted 1 times

**JasMozai** 2 years, 5 months ago

No definition of Global Services Layer, it should be cloud services layer Snowflake stores metadata about all rows stored in a micro-partition, not columns

upvoted 3 times

**lol632** 2 years, 11 months ago

https://www.snowflake.com/wp-content/uploads/2014/10/A-Detailed-View-Inside-Snowflake.pdf It clearly says the service layer saves the metadata about the stored data in snowflake. Page 7, right column.

upvoted 4 times

**oscarglob** 2 years, 11 months ago

"Snowflake stores metadata about all rows stored in a micro-partition, including: - The range of values for each of the columns in the micro-partition. - The number of distinct values. - Additional properties used for both optimization and efficient query processing." Seems to me that this is done by the Cloud Services layer, and this "Global Services layer" looks like another term for it: https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions#what-are-micro-partitions

upvoted 3 times

**dpexam2023** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions

upvoted 1 times

**Marge23** 3 years, 1 month ago

B . FALSE Cloud Services If Snowflake's computing layer is the brawn of our solution, our cloud services layer is the brain. Snowflake's cloud services manage client sessions, metadata, transactions, query planning, security/governance and many other services. It's also a highly scalable tier thanks to the near-limitless computing resources in the cloud. Snowflake supports thousands of customer accounts and hundreds of millions of queries each and every day.

upvoted 2 times

**MultiCloudIronMan** 3 years, 3 months ago

There is no Global Services Layer in snowflake. https://www.bing.com/search?q=snowflake+global+services+layer&qs=NM&pq=what+is+global+services+layer+in+snowflake&sc=6-42&cvid=9660BE0F9FA3415AB5F2EF9404F1C104&FORM=QBRE&sp=1&ghc=2&lq=0

upvoted 2 times

**Leinho** 3 years, 4 months ago

B is correct, question is about micro-partitions, documentation explains it keeps metadata about rows, not columns, so is false.

upvoted 3 times

**Leinho** 3 years, 4 months ago

B is correct, question is about micro-partitions, documentation explains it keeps metadata about rows, not columns, so is false.metadata comes from the micro partition data only right. metadata is count(*), min and max values in each micro partition, distinct values. min and max of each micro partition is what it decides pruning right?

upvoted 3 times

**Rajivnb** 3 years ago

metadata comes from the micro partition data only right. metadata is count(*), min and max values in each micro partition, distinct values. min and max of each micro partition is what it decides pruning right?

upvoted 1 times

**tochengiri** 3 years, 5 months ago

https://www.snowflake.com/product/architecture/ Snowflake is a single, integrated platform delivered as-a-service. It features storage, compute, and global services layers that are physically separated but logically integrated. Data workloads scale independently from one another, making it an ideal platform for data warehousing, data lakes, data engineering, data science, modern data sharing, and developing data applications.

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**EmiB** 3 years, 7 months ago

Snowflake is a single, integrated platform delivered as-a-service. It features storage, compute, and GLOBAL SERVICES LAYERS
https://www.snowflake.com/product/architecture/

upvoted 1 times

Question #93 *Topic 1*

True or False: It is best practice to define a clustering key on every table.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

NO,It is recommended to use on very large tables(> 1 TB)

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct - it used for very large tables

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 2 times

**SteveCrown** 3 years, 6 months ago

Answer: B It is not always best practice in Snowflake to define a clustering key on every table. Clustering keys are used to group related rows physically together in a table, which can improve query performance by minimizing the amount of data that needs to be scanned or skipped during query execution. However, clustering keys should be chosen carefully based on the access patterns of the table, and the type of queries that will be run against it. Defining a clustering key on every table may not always be necessary or beneficial. For example, in some cases, tables may be small or may not have well-defined access patterns, making clustering unnecessary. In other cases, tables may be loaded in a specific order or partitioned in a certain way that obviates the need for clustering. Therefore, it is not a best practice to define a clustering key on every table in Snowflake, but rather to carefully consider the access patterns and usage of each table, and to choose clustering keys judiciously based on those factors.

upvoted 4 times

**Tapasgup007** 3 years, 10 months ago

B is correct. It is not a best practice

upvoted 1 times

**moon2351** 4 years, 10 months ago

Answer is correct.

upvoted 1 times

Question #94 *Topic 1*

Which of the following statements is true of Snowflake?

A. It was built specifically for the cloud

B. It was built as an on-premises solution and then ported to the cloud

C. It was designed as a hybrid database to allow customers to store data either on premises or in the cloud

D. It was built for Hadoop architecture

E. It's based on an Oracle Architecture

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Harshh_zz** Most Recent 1 year, 2 months ago

Build for the cloud

upvoted 1 times

**francisss** 1 year, 8 months ago

Snowflake is provided as a self-managed service that runs completely on cloud infrastructure. This means that all three layers of Snowflake's architecture (storage, compute, and cloud services) are deployed and managed entirely on a selected cloud platform. A Snowflake account can be hosted on any of the following cloud platforms: Amazon Web Services (AWS) Google Cloud Platform (GCP) Microsoft Azure (Azure) source: https://docs.snowflake.com/en/user-guide/intro-cloud-platforms

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A.It was built specifically for the cloud

upvoted 1 times

**Colourseun** 2 years, 12 months ago

A is right Snowflake is a Cloud only solution

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**lhstudylh** 4 years, 8 months ago

Correct answer

upvoted 1 times

Question #95 *Topic 1*

What is the minimum Snowflake edition that provides multi-cluster warehouses and up to 90 days of Time Travel?

A. Standard

B. Premier

C. Enterprise

D. Business Critical Edition

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

C.Enterprise Edition

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

C is correct

upvoted 1 times

**cgapperi** 3 years, 8 months ago

There is no Premier edition, so not B. :) Correct answer is C

upvoted 1 times

**fahfouhi94** 3 years, 8 months ago

Entreprise is the min snow edition

upvoted 1 times

**kjdfsafc87456** 4 years, 2 months ago

The answer is correct as B is not under the snowflake options; and extended time travel up to 90 days is available for Enterprise, Business Critical and VPS

upvoted 1 times

**Rajokkiyam** 4 years, 8 months ago

Correct Answer - Enterprise. Standard Offer Timetravel up to 1 day. Enterprise + edition offer up to 90 days.

upvoted 3 times

Question #96 *Topic 1*

How many shares can be consumed by a single Data Consumer?

A. 1

B. 10

C. 100, but can be increased by contacting support

D. Unlimited

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**56d43d6** Most Recent 1 year, 10 months ago

D: https://docs.snowflake.com/en/user-guide/data-sharing-intro#about-consumers

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

D.is correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

D is correct

upvoted 1 times

**AravindhTN** 3 years, 10 months ago

Correct

upvoted 1 times

**learnazureportal** 4 years, 5 months ago

Answer is correct.

upvoted 2 times

Question #97 *Topic 1*

What is the lowest Snowflake edition that offers Time Travel up to 90 days?

A. Standard Edition

B. Premier Edition

C. Enterprise Edition

D. Business Critical Edition

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

C. Enterprise Edition

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

C is correct

upvoted 1 times

**EmiB** 3 years, 7 months ago

C. Enterprise Edition

upvoted 1 times

**addixion** 4 years, 7 months ago

Only Standard has 1 day

upvoted 1 times

Question #98 *Topic 1*

Which of the following statements are true about Schemas in Snowflake? (Choose two.)

A. A Schema may contain one or more Databases

B. A Database may contain one or more Schemas

C. A Schema is a logical grouping of Database Objects

D. Each Schema is contained within a Warehouse

**Suggested Answer:** *B, C*

*Community vote distribution*

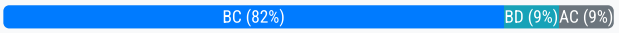


**BrunoRYamada** Most Recent 1 year, 7 months ago

A & C B is wrong because DBs don't contain schema, it is the other way around

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B and C are Correct Options

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

BC is correct

upvoted 2 times

**Raju039** 3 years, 8 months ago

BC is correct . Warehouse is compute, it doesn't contain schema within

upvoted 2 times

**c0d3g** 3 years, 8 months ago

Warehouse, here means compute and not traditional warehouse.

upvoted 1 times

**rajivgupta780** 3 years, 9 months ago

BC IS CORRECT

upvoted 1 times

**AlexbDku** 3 years, 10 months ago

Answer is BC

upvoted 3 times

**AlexbDku** 3 years, 10 months ago

Answer is BD

upvoted 1 times

**learnazureportal** 4 years, 6 months ago

Answer is correct.

upvoted 2 times

**grabc** 4 years, 9 months ago

How about BD?

upvoted 1 times

**grabc** 4 years, 9 months ago

How about BD?Warehouse - In Snowflake world, corresponds to compute resource.

upvoted 1 times

**Rajokkiyam** Highly Voted 4 years, 8 months ago

Warehouse - In Snowflake world, corresponds to compute resource.

upvoted 5 times

Question #99 *Topic 1*

True or False: You can resize a Virtual Warehouse while queries are running.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*

A (88%) B (12%)

**jtm22** Highly Voted 3 years, 9 months ago

A warehouse can be resized up or down at any time, including while it is running and processing statements.

upvoted 5 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A Is correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 2 times

**ka2456** 3 years, 9 months ago

False. You cannot resize a Virtual Warehouse in Snowflake while queries are running. You must first pause or terminate any running queries before resizing the Virtual Warehouse.Not correct, you can resize Warehouse anytime, it just won't use new resources for the already running queries.no you can resize it but it won't affect the current query. so changes apply for next ones

upvoted 1 times

**ka2456** 3 years, 9 months ago

Selected Answer: B

False. You cannot resize a Virtual Warehouse in Snowflake while queries are running. You must first pause or terminate any running queries before resizing the Virtual Warehouse.

upvoted 1 times

**SalihAlic** 3 years, 9 months ago

Not correct, you can resize Warehouse anytime, it just won't use new resources for the already running queries.

upvoted 4 times

**Hopefan** 3 years, 5 months ago

no you can resize it but it won't affect the current query. so changes apply for next ones

upvoted 3 times

**learnazureportal** 4 years, 6 months ago

Answer is correct

upvoted 3 times

Question #100 *Topic 1*

What is the most granular object that the Time Travel retention period can be defined on?

A. Account

B. Database

C. Schema

D. Table

**Suggested Answer:** *D*

*Community vote distribution*

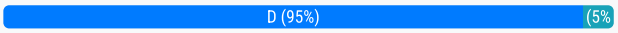


**blazingliverbird** Highly Voted 4 years, 10 months ago

The time travel data retention can be overwritten at the table level "When creating a table, schema, or database, the account default can be overridden using the DATA_RETENTION_TIME_IN_DAYS parameter in the command."

upvoted 14 times

**Sid460545** Highly Voted 4 years, 9 months ago

Granular is Table. It is the most granular object not account. Account--database --schema-- table(most granular)

upvoted 8 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

D.Table

upvoted 1 times

**6c79d6f** 2 years, 10 months ago

The Time Travel retention period in Snowflake specifies the number of days for which historical data is preserved and, therefore, Time Travel operations (SELECT, CREATE … CLONE, UNDROP) can be performed on the data 1. The standard retention period is 1 day (24 hours) and is automatically enabled for all Snowflake accounts 1. For Snowflake Enterprise Edition (and higher), the retention period can be set to any value from 0 up to 90 days for permanent databases, schemas, and tables 12. Therefore, the most granular object that the Time Travel retention period can be defined on is a table 1. I hope this helps! https://docs.snowflake.com/en/user-guide/data-time-travel

upvoted 1 times

**6c79d6f** 2 years, 11 months ago

High granularity level refers to a high level of detail, vice-versa low granularity level refers to a low level of detail. Practically speaking, the more subdividable and specific a data is, the more granular it is considered to be. Thus, “granularity” and “level of detail” of data are the same thing. In this case the most granular level is asked and that is a syntactic issue as most is defined by the perspective. Highest or Lowest level So both A and D could be correct. However I think that Most is meant als the highest level of detail and that would be D.

upvoted 1 times

**JFalcon** 3 years ago

D is correct

upvoted 1 times

**ifmateus** 3 years, 2 months ago

Most granular means with the highest level of detail, where do we obtain that? Of course in the table, the account contains a low level of detail, answer is D.

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**EmiB** 3 years, 7 months ago

D. Table [ALTER TABLE t SET DATA_RETENTION_TIME_IN_DAYS = 0;]

upvoted 1 times

**TWL009** 3 years, 8 months ago

Dear Admin most granular means lowest now here answer would be table. Pls correct the answers

upvoted 1 times

**KarBiswa** 3 years, 9 months ago

Table is most granular

upvoted 2 times

**ajay_1233456** 3 years, 9 months ago

D is correct answer, at the lowest granular level object time travel can be configured.

upvoted 1 times

**baig123** 3 years, 10 months ago

D is correct answer, at the lowest granular level object time travel can be configured.

upvoted 1 times

**Rajokkiyam** 4 years, 8 months ago

Answer: Table

upvoted 2 times

**JustVicky** 4 years, 9 months ago

To specify the data retention period for Time Travel: The DATA_RETENTION_TIME_IN_DAYS object parameter can be used by users with the ACCOUNTADMIN role to set the default retention period for your account. The same parameter can be used to explicitly override the default when creating a database, schema, and individual table. The data retention period for a database, schema, or table can be changed at any time.

upvoted 3 times

Question #101 *Topic 1*

Which of the following statements is true of Snowflake micro-partitioning?

A. Micro-partitioning has been known to introduce data skew

B. Micro-partitioning: requires a partitioning schema to be defined up front

C. Micro-partitioning is transparently completed using the ordering that occurs when the data is inserted/loaded

D. Micro-partitioning can be disabled within a Snowflake account

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**ArunRav** Most Recent 1 year, 3 months ago

https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions#what-are-micro-partitions Tables are transparently partitioned using the ordering of the data as it is inserted/loaded.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

C is Correct

upvoted 1 times

**Marge23** 3 years, 1 month ago

correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

C is correct

upvoted 1 times

**Bagvi** 3 years, 7 months ago

Option C is correct

upvoted 1 times

**c0d3g** 3 years, 8 months ago

Option C

upvoted 1 times

**DK69** 4 years, 5 months ago

C is correct

upvoted 1 times

Question #102 *Topic 1*

True or False: Snowflake bills for a minimum of five minutes each time a Virtual Warehouse is started.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**arpit_dataguy** Highly Voted 4 years, 9 months ago

It bills for minimum 1 minute.

upvoted 8 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

No,It bills for minimum of 60 sec,B.False

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**Bagvi** 3 years, 7 months ago

It dos't take 5 inutes

upvoted 1 times

**AlexbDku** 3 years, 10 months ago

Because Snowflake utilizes per-second billing (with a 60-second minimum each time the warehouse starts), warehouses are billed only for the credits they actually consume when they are actively working.

upvoted 2 times

Question #103 *Topic 1*

When scaling up Virtual Warehouses by increasing Virtual Warehouse t-shirt size, you are primarily scaling for improved:

A. Concurrency

B. Performance

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Surbin** Highly Voted 3 years, 6 months ago

Scaling up for performance and scaling out for concurrency

upvoted 8 times

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

Correct

upvoted 1 times

**Leinho** 2 years, 5 months ago

What is a t-shirt size?

upvoted 2 times

**Leinho** 2 years, 5 months ago

What is a t-shirt size?When we create a warehouse, we need to assign the number of servers it will use. So this varying number of servers are represented in form of your regular T-Shirt sizes, like XL,L,XXL. 1 Server => XS(extra small) 2 Server => S(Small) 4 Server => M(Medium) 6 Server = > L(Large) 8 Server => XL (Extra large) it goes on, XXL,XXXL Notice the step count is +2 (That's snowflake defined)Correction:- The step size is "x2" not "+2" (Multiplied by a factor of 2) So it's, 2,4,8,16,32,64,128

upvoted 2 times

**SouravRvd** 2 years, 4 months ago

When we create a warehouse, we need to assign the number of servers it will use. So this varying number of servers are represented in form of your regular T-Shirt sizes, like XL,L,XXL. 1 Server => XS(extra small) 2 Server => S(Small) 4 Server => M(Medium) 6 Server = > L(Large) 8 Server => XL (Extra large) it goes on, XXL,XXXL Notice the step count is +2 (That's snowflake defined)Correction:- The step size is "x2" not "+2" (Multiplied by a factor of 2) So it's, 2,4,8,16,32,64,128

upvoted 1 times

**SouravRvd** 2 years, 4 months ago

When we create a warehouse, we need to assign the number of servers it will use. So this varying number of servers are represented in form of your regular T-Shirt sizes, like XL,L,XXL. 1 Server => XS(extra small) 2 Server => S(Small) 4 Server => M(Medium) 6 Server = > L(Large) 8 Server => XL (Extra large) it goes on, XXL,XXXL Notice the step count is +2 (That's snowflake defined)

upvoted 1 times

**SouravRvd** 2 years, 4 months ago

Correction:- The step size is "x2" not "+2" (Multiplied by a factor of 2) So it's, 2,4,8,16,32,64,128

upvoted 1 times

**Mayuri917233** 2 years, 5 months ago

B is correct

upvoted 1 times

**Bagvi** 2 years, 7 months ago

Option B

upvoted 1 times

**jobbadahut** 3 years, 2 months ago

Not sure on this one

upvoted 1 times

**aqmdla2002** 3 years, 3 months ago

B is correct

upvoted 1 times

Question #104 *Topic 1*

As a best practice, clustering keys should only be defined on tables of which minimum size?

A. Multi-Kilobyte (KB) Range

B. Multi-Megabyte (MB) Range

C. Multi-Gigabyte (GB) Range

D. Multi-Terabyte (TB) Range

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**PrasantSadangi** Most Recent 1 year, 7 months ago

Considerations for Choosing Clustering for a Table Whether you want faster response times or lower overall costs, clustering is best for a table that meets all of the following criteria: The table contains a large number of micro-partitions. Typically, this means that the table contains multiple terabytes (TB) of data. The queries can take advantage of clustering. Typically, this means that one or both of the following are true: The queries are selective. In other words, the queries need to read only a small percentage of rows (and thus usually a small percentage of micro-partitions) in the table. The queries sort the data. (For example, the query contains an ORDER BY clause on the table.) https://docs.snowflake.com/en/user-guide/tables-clustering-keys

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

D.is correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

D is correct

upvoted 1 times

**Bagvi** 3 years, 7 months ago

Multi TB range

upvoted 2 times

**abbcia** 3 years, 7 months ago

saw it in their core course

upvoted 1 times

**datamech001** 4 years ago

D is correct Considerations for Choosing Clustering for a Table Whether you want faster response times or lower overall costs, clustering is best for a table that meets all of the following criteria: The table contains a large number of micro-partitions. Typically, this means that the table contains multiple terabytes (TB) of data.

upvoted 3 times

**DK69** 4 years, 5 months ago

D is correct

upvoted 1 times

Question #105 *Topic 1*

How a Snowpipe charges calculated?

A. Per-second/per Warehouse size

B. Per-second/per-core granularity

C. Number of Pipes in account

D. Total storage bucket size

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**lhstudylh** Highly Voted 4 years, 8 months ago

Correct answer "Snowflake tracks the resource consumption of loads for all pipes in an account, with per-second/per-core granularity, as Snowpipe actively queues and processes data files. Per-core refers to the physical CPU cores in a compute server. The utilization recorded is then translated into familiar Snowflake credits, which are listed on the bill for your account."

upvoted 11 times

**PrasantSadangi** Most Recent 1 year, 7 months ago

https://docs.snowflake.com/en/user-guide/data-load-snowpipe-billing#resource-consumption-and-management-overhead Resource consumption and management overhead Snowflake tracks the resource consumption of loads for all pipes in an account, with per-second/per-core granularity, as Snowpipe actively queues and processes data files. Per-core refers to the physical CPU cores in a compute server. The utilization recorded is then translated into familiar Snowflake credits, which are listed on the bill for your account.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B is correct

upvoted 1 times

**Stevec80** 2 years, 5 months ago

https://docs.snowflake.com/en/user-guide/data-load-snowpipe-billing#resource-consumption-and-management-overhead

upvoted 1 times

**lofibeats** 3 years, 2 months ago

answer is correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**Bagvi** 3 years, 7 months ago

Calculated per second per core granularity

upvoted 1 times

**greenlever** 4 years ago

Snowpipe's serverless compute model, users can initiate any size load without managing a virtual warehouse

upvoted 2 times

**DK69** 4 years, 5 months ago

Answer is B. Snowpipe is serverless

upvoted 1 times

**Sly007** 4 years, 6 months ago

Answer is A. each warehouse size has different credit charge when used

upvoted 1 times

Question #106 *Topic 1*

True or False: A Snowflake account is charged for data stored in both Internal and External Stages.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*




**visionary4ever** Most Recent 1 year, 2 months ago

Internal stages: Data is stored inside Snowflake-managed cloud storage → Snowflake charges you for this storage. External stages: Data lives in your own cloud storage (AWS S3, Azure Blob, GCS). Snowflake still bills you for metadata operations and access, but the cloud provider separately charges for the storage itself. So, from the perspective of a Snowflake account: yes, you incur charges for both — either directly (internal) or indirectly (external).

upvoted 1 times

**7f9e18c** 1 year, 8 months ago

The external stage is not part of Snowflake, so Snowflake does not store or manage the stage

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

Snowflake doesn't charge for both,B.False

upvoted 1 times

**Amitsnowflake** 3 years, 1 month ago

For an external stage, cloud provider charge you bill but for internal stage snowflake will Charge. Since question asked for both then ans is false, snowflake doesn't charge for both

upvoted 4 times

**lofibeats** 3 years, 2 months ago

Snowflake supports both internal stages and external stages. Data files staged in Snowflake internal stages are not subject to the additional costs associated with Time Travel and Fail-safe, but they do incur standard data storage costs. https://docs.snowflake.com/en/user-guide/tables-storage-considerations

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**Bagvi** 3 years, 7 months ago

Only in external store

upvoted 1 times

**c0d3g** 3 years, 8 months ago

The customer manages external stage

upvoted 1 times

**rajivgupta780** 3 years, 9 months ago

EXTERNAL STAGE BILLING IS OUTSIDE SNOWFLAKE ACCOUNT BILLING HENCE, B IS CORRECT

upvoted 1 times

**vetaal** 4 years ago

False, the question says Account is charged for Data stored - but account is charged for overheads related to autorefresh checks.

upvoted 1 times

**aqmdla2002** 4 years, 3 months ago

https://docs.snowflake.com/en/user-guide/tables-external-intro.html#billing-for-external-tables For external, there are some overheads that will be billed

upvoted 1 times

**distopic** 4 years, 4 months ago

External data is not billed

upvoted 1 times

**Sid460545** 4 years, 9 months ago

Not for External Stage.. So False

upvoted 2 times

**arpit_dataguy** 4 years, 9 months ago

False. There are no storage charges for external stages billed through Snowflake. We need to pay the external vendor directly and not to the snowflake.

upvoted 4 times

**JustVicky** 4 years, 9 months ago

False. There are no storage charges for external stages billed through Snowflake.

upvoted 3 times

Question #107 *Topic 1*

True or False: When active, a Pipe requires a dedicated Virtual Warehouse to execute.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**lhstudylh** Highly Voted 4 years, 8 months ago

Correct Answer Bulk-loading: Requires virtual warehouse to use COPY INTO command. Snowpipe (continuous loading): Uses Snowflake-supplied compute resources and charges based on how much of that is used.

upvoted 9 times

**someshVashistha** Most Recent 1 year, 10 months ago

Snowpipe is serverless service

upvoted 3 times

**Mallikharjuna452** 2 years, 3 months ago

B Is Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**Bagvi** 3 years, 7 months ago

Pipe not required

upvoted 1 times

**kjdfsafc87456** 4 years, 2 months ago

why do we need a VW to run a serverless service?

upvoted 3 times

Question #108 *Topic 1*

True or False: Snowflake supports federated authentication in all editions.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*

A (78%) B (22%)

**56d43d6** Most Recent 1 year, 10 months ago

A true, federated auth is included in Standard and everything above it. https://docs.snowflake.com/en/user-guide/intro-editions#security-governance-and-data-protection

upvoted 4 times

**vhvhkkk** 1 year, 11 months ago

Federated authentication is available with organisations feature. and organisations is available in enterprise and above.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

Yes,It support in all editions

upvoted 1 times

**lol632** 2 years, 11 months ago

Yes it's A

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**Bagvi** 3 years, 7 months ago

Yes, it supports

upvoted 1 times

**greenlever** 4 years ago

Support for user SSO (single sign-on) through federated authentication - All Edition

upvoted 3 times

**seaun** 4 years ago

A is Correct

upvoted 1 times

**Naveen90** 4 years, 1 month ago

Answer is A - True

upvoted 1 times

**kjdfsafc87456** 4 years, 2 months ago

Selected Answer: B

Get the answer by searching federated authentication from https://docs.snowflake.com/en/user-guide/admin-security.html

upvoted 1 times

**kjdfsafc87456** 4 years, 2 months ago

Get the answer by searching federated authentication from https://docs.snowflake.com/en/user-guide/admin-security.htmlThen why did you select B? It should A

upvoted 1 times

**greenlever** 4 years ago

Then why did you select B? It should A

upvoted 1 times

**JustVicky** 4 years, 9 months ago

Support for user SSO through federated authentication

upvoted 2 times

Question #109 *Topic 1*

True or False: When a new Snowflake object is created, it is automatically owned by the user who created it.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*




**Surbin** Highly Voted 4 years, 7 months ago

Owner is always the role, not the user.

upvoted 24 times

**c0d3g** Highly Voted 3 years, 8 months ago

Each object has an owner, who can in turn grant access to that object. To own an object means that a role has the OWNERSHIP privilege on the object. Each securable object is owned by a single role, which by default is the role used to create the object. When this role is assigned to users, they effectively have shared control over the object. In a regular schema, the owner role has all privileges on the object by default, including the ability to grant or revoke privileges on the object to other roles. https://docs.snowflake.com/en/user-guide/security-access-control-overview

upvoted 6 times

**jeny12138** Most Recent 1 year, 7 months ago

True. In Snowflake, when a new object (like a table, view, schema, etc.) is created, it is automatically owned by the user (or more precisely, the role) that created it. Ownership includes full control over the object, including the ability to: Grant privileges to other roles, Modify or drop the object, Transfer ownership.

upvoted 2 times

**56d43d6** 1 year, 10 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-overview#securable-objects

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

Yes it is True

upvoted 1 times

**KnightVictor** 3 years, 5 months ago

owned by role, and not by user

upvoted 3 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 2 times

**isolated_sniper** 3 years, 6 months ago

Any objects created after the command is issued are owned by the role in use when the object is created. https://docs.snowflake.com/en/sql-reference/sql/grant-ownership#

upvoted 2 times

**Bagvi** 3 years, 7 months ago

Not own by the users

upvoted 3 times

**Raju039** 3 years, 8 months ago

owned by the role of user who created it

upvoted 1 times

**alloshz11** 3 years, 8 months ago

To own an object means that a role has the OWNERSHIP privilege on the object. Each securable object is owned by a single role, which by default is the role used to create the object. When this role is assigned to users, they effectively have shared control over the object. In a regular schema, the owner role has all privileges on the object by default, including the ability to grant or revoke privileges on the object to other roles. In addition, ownership can be transferred from one role to another https://docs.snowflake.com/en/user-guide/security-access-control-overview#securable-objects

upvoted 2 times

**MultiCloudIronMan** 3 years, 8 months ago

This is true its called DAC - Discretional Access Control.

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

When an object is created, its ownership is set to the currently active primary role.

upvoted 1 times

**BigDataBB** 3 years, 10 months ago

It's an ambiguous question, if was ... automatically and "exclusively" owned by the user who created it. Then is sure B but in this way, as long as that user has that role, he has the ownersiph of the objet can alter it and even drop it.

upvoted 2 times

**Bala1802** 3 years, 10 months ago

Can we assume a user as a credential you log in with & role is what you privilege to do Also can we assume an object created by a role can be accessible to all users under the same role?

upvoted 1 times

**SV1122** 3 years, 10 months ago

**Selected Answer: A**

Discretionary Access Control (DAC): Each object has an owner, who can in turn grant access to that object.

Role-based Access Control (RBAC): Access privileges are assigned to roles, which are in turn assigned to users.

https://docs.snowflake.com/en/user-guide/security-access-control-overview.html

upvoted 2 times

**SV1122** 3 years, 10 months ago

Discretionary Access Control (DAC): Each object has an owner, who can in turn grant access to that object. Role-based Access Control (RBAC): Access privileges are assigned to roles, which are in turn assigned to users. https://docs.snowflake.com/en/user-guide/security-access-control-overview.htmlEach object has an owner, who can in turn grant access to that object. To own an object means that a role has the OWNERSHIP privilege on the object. Each securable object is owned by a single role, which by default is the role used to create the object. When this role is assigned to users, they effectively have shared control over the object. In a regular schema, the owner role has all privileges on the object by default, including the ability to grant or revoke privileges on the object to other roles. https://docs.snowflake.com/en/user-guide/security-access-control-overview

upvoted 2 times

**baig123** 3 years, 10 months ago

Answer is A, it is DAC ownership assigned to role under which object is created.

upvoted 1 times

Question #110 *Topic 1*

True or False: A Virtual Warehouse consumes Snowflake credits even when inactive.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*




**alexhparreira** Most Recent 5 months ago

Why it is False: A Virtual Warehouse only consumes Snowflake credits when it is actively running and executing queries or data loading operations. When a warehouse is in an Inactive (Suspended) state, it releases its underlying cloud compute resources and completely stops consuming compute credits.

upvoted 1 times

**niho1993** 1 year, 1 month ago

By default, all warehouses have the auto-suspend setting enabled, which means a warehouse automatically shuts down when it is inactive for a defined period of time. A suspended warehouse does not consume credits, so the warehouse only incurs cost when it is processing a workload.

upvoted 1 times

**Santa9os** 1 year, 2 months ago

A warehouse continues to consume credits if it is still technically running, even if it's inactive (not executing queries). It only stops consuming credits when it is fully suspended. Running → Inactive → Auto-suspend → Suspended

upvoted 2 times

**PrasantSadangi** 1 year, 7 months ago

Inactive is not equal to suspended. Until the inactivity timeout not reached & warehouse is not suspended the warehouse will consume credit.

upvoted 3 times

**Siif97** 1 year, 9 months ago

In Snowflake, "inactive" typically means that the Virtual Warehouse is suspended. When a warehouse is suspended, it does not consume any credits

upvoted 1 times

**cgm_e** 1 year, 9 months ago

Suspended ≠ inactive. A Warehouse will be suspended if inactive for a predetermined period of time between 1 minute and 365 days.

upvoted 3 times

**rubiteb** 1 year, 11 months ago

Inactive - could mean idle but not suspended. This still consume credits.

upvoted 1 times

**Srisankar** 2 years, 2 months ago

False, As Inactive or suspended warehouse do not incur charges.

upvoted 1 times

**capt2101akash** 2 years, 3 months ago

Suspend i.e inative

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B.False,When it is inactive or suspended, it does not consume credits.

upvoted 1 times

**PrashantGupta1616** 2 years, 4 months ago

Yes, Suspend is not equal to Inactive Charges will keep incur till get suspended/off

upvoted 3 times

**MBrat** 2 years, 5 months ago

Inactive = suspended according to https://docs.snowflake.com/en/user-guide/warehouses-tasks#starting-or-resuming-a-warehouse "A suspended (i.e. inactive) warehouse can be resumed through the web interface or using SQL"

upvoted 4 times

**Stevec80** 2 years, 5 months ago

Snowflake charges only when the warehouse is running. No cost when it's inactive (no running for 'tot' period) or suspended

upvoted 1 times

**abhijit2683** 2 years, 7 months ago

A is correct

upvoted 2 times

**rogerek077** 2 years, 7 months ago

Check this document: https://docs.snowflake.com/en/user-guide/warehouses-overview#auto-suspension-and-auto-resumption "By default, auto-suspend is enabled. Snowflake automatically suspends the warehouse if it is inactive for the specified period of time." At that point Snowflake makes clear distinction between "active/inactive" and "suspended". And considering this article: https://docs.snowflake.com/en/user-guide/cost-understanding-compute#virtual-warehouse-credit-usage The only conclusion is that the answer to the question is A - Virtual Warehouse consumes credits even if it's inactive, as long as it's not suspended manually or by the auto-suspend functionality.

upvoted 3 times

**junaid2107** 2 years, 11 months ago

is inactive=suspended?

upvoted 1 times

**sfdumps** 2 years, 10 months ago

WAREHOUSE IS BILLED UNLESS IT IS NOT SUSPENDED

upvoted 3 times

**junaid2107** 2 years, 11 months ago

is inactive=suspended?These two are not the same. Check this article: https://docs.snowflake.com/en/user-guide/warehouses-overview#auto-suspension-and-auto-resumption "By default, auto-suspend is enabled. Snowflake automatically suspends the warehouse if it is inactive for the specified period of time."

upvoted 1 times

**iamPhD** 3 years ago

**Selected Answer: A**

Warehouse is billed even if it is inactive unless suspended

upvoted 4 times

**rogerek077** 2 years, 7 months ago

These two are not the same. Check this article: https://docs.snowflake.com/en/user-guide/warehouses-overview#auto-suspension-and-auto-resumption "By default, auto-suspend is enabled. Snowflake automatically suspends the warehouse if it is inactive for the specified period of time."

upvoted 3 times

Question #111 *Topic 1*

True or False: During data unloading, only JSON and CSV files can be compressed.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*




**lakediver** Highly Voted 2 years, 6 months ago

Along with JSON and CSV - Parquet files are also compressed while unloading https://docs.snowflake.com/en/sql-reference/sql/copy-into-location

upvoted 7 times

**isolated_sniper** Highly Voted 3 years, 6 months ago

During data loading, compression can be enabled for CSV, JSON PARQUET, AVRO, XML type files. ref: https://docs.snowflake.com/en/sql-reference/sql/create-file-formatthis question is about unloading. avro and xml are not relevant for unloading.question is on data 'unloading'

upvoted 5 times

**isolated_sniper** 3 years, 6 months ago

During data loading, compression can be enabled for CSV, JSON PARQUET, AVRO, XML type files. ref: https://docs.snowflake.com/en/sql-reference/sql/create-file-format

upvoted 5 times

**56d43d6** 1 year, 10 months ago

this question is about unloading. avro and xml are not relevant for unloading.

upvoted 3 times

**subhataanu** 2 years, 5 months ago

question is on data 'unloading'

upvoted 2 times

**f3b379a** Most Recent 9 months ago

during unloading data can be compressed

upvoted 1 times

**Shushmita** 1 year, 8 months ago

A, only csv and json can be compressed during unloading

upvoted 1 times

**56d43d6** 1 year, 10 months ago

B: https://docs.snowflake.com/en/sql-reference/sql/copy-into-location#format-type-options-formattypeoptions includes csv, json and parquet

upvoted 1 times

**iusechegg** 1 year, 10 months ago

https://docs.snowflake.com/en/user-guide/data-unload-prepare Parquet is also an option

upvoted 1 times

**sg242** 1 year, 11 months ago

The question talks about unloading that is GET command. It support compression for csv and json only.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B.False

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. False

upvoted 2 times

**Aaron_w_mendelson** 3 years ago

B "By default, all unloaded data files are compressed using gzip, unless compression is explicitly disabled or one of the other supported compression methods is explicitly specified." https://docs.snowflake.com/en/user-guide/intro-summary-unloading

upvoted 3 times

**Thaiya** 3 years, 1 month ago

only JSON and CSV formats can be unloaded into. Parquet AVro etc only for loading.

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**RajuNaik** 3 years, 6 months ago

Data can be unloaded into csv,json and Parquet. so answer is False

upvoted 3 times

**alanhuangjob** 3 years, 7 months ago

CSV for structured data; JSON and Parquet for semi-structured data.

upvoted 2 times

**Sk1020** 4 years ago

Unloading supports structured formatting CSV format and unstructured JSON and Parquet formats only. Hence FALSE is correct answer.

upvoted 2 times

**Sid460545** 4 years, 9 months ago

Also Parquet file.. So false

upvoted 2 times

**arpit_dataguy** 4 years, 9 months ago

csv, json & parquet can be compressed

upvoted 3 times

Question #112 *Topic 1*

Which of the following are options when creating a Virtual Warehouse? (Choose two.)

A. Auto-suspend

B. Auto-resume

C. Local SSD size

D. User count

**Suggested Answer:** *A, B*

*Community vote distribution*

AB (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A and B are correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

AB are correct

upvoted 1 times

**fahfouhi94** 3 years, 8 months ago

AB ===> correct answer

upvoted 1 times

**DK69** 4 years, 5 months ago

Answer is correct

upvoted 1 times

Question #113 *Topic 1*

Which formats are supported for unloading data from Snowflake? (Choose two.)

A. Delimited (CSV, TSV, etc.)

B. Avro

C. JSON

D. ORC

**Suggested Answer:** *A, C*

*Community vote distribution*

AC (100%)

**AbhaySAP** Most Recent 2 years, 4 months ago

AC are correct

upvoted 1 times

**Mayuri917233** 2 years, 5 months ago

AC are correct

upvoted 1 times

**SV1122** 2 years, 10 months ago

Was on exam Dec 31st 2022always Dec 31st 2022. Be creative please. Not many people want to write an exam on new year's eve.

upvoted 1 times

**SV1122** 2 years, 10 months ago

Was on exam Dec 31st 2022

upvoted 1 times

**jw1806** Highly Voted 2 years, 7 months ago

always Dec 31st 2022. Be creative please. Not many people want to write an exam on new year's eve.

upvoted 13 times

**SV1122** 2 years, 10 months ago

Supported file formats Delimited files (CSV, TSV, etc.) JSON Parquet

upvoted 2 times

**greenlever** 3 years ago

correct.

upvoted 1 times

**kjdfsafc87456** 3 years, 2 months ago

The answer is correct as snow supported structured data file format as type of delimited such as CSV etc and etc while semi-structured type of Json and parquet

upvoted 1 times

**Sid460545** 3 years, 9 months ago

Correct, as per options given A and C

upvoted 1 times

**arpit_dataguy** 3 years, 9 months ago

csv, json & parquet

upvoted 2 times

**JustVicky** 3 years, 9 months ago

AC is correct. https://docs.snowflake.com/en/user-guide/intro-summary-unloading.html -- Output data file details table

upvoted 1 times

**blazingliverbird** 3 years, 10 months ago

All the options given are supported.

upvoted 3 times

**blazingliverbird** 3 years, 10 months ago

All the options given are supported.Not all supported.. https://docs.snowflake.com/en/user-guide/data-unload-prepare.html

upvoted 3 times

**arunshai** 3 years, 9 months ago

Not all supported.. https://docs.snowflake.com/en/user-guide/data-unload-prepare.html

upvoted 2 times

Question #114 *Topic 1*

True or False: Data Providers can share data with only the Data Consumer.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*




**greenlever** Highly Voted 4 years ago

Selected Answer: B

Data Providers can share data with reader also. So it is False(B)

upvoted 13 times

**greenlever** Highly Voted 4 years ago

Data Providers can share data with reader also. So it is False(B)If the question means to distinguish between a reader account set up by a Provider account, and an external Snowflake account which consumes a data share from the provider, then I think your answer makes sense. However, I'm not so sure that a reader and a consumer are considered different things.I am sorry for the spam: but here in the snowflake documentation it says a reader is still considered a consumer: https://docs.snowflake.com/en/user-guide/data-sharing-reader-create So I think your answer here is incorrect.

upvoted 13 times

**AkexS** 3 years, 3 months ago

If the question means to distinguish between a reader account set up by a Provider account, and an external Snowflake account which consumes a data share from the provider, then I think your answer makes sense. However, I'm not so sure that a reader and a consumer are considered different things.

upvoted 1 times

**AkexS** 3 years, 3 months ago

If the question means to distinguish between a reader account set up by a Provider account, and an external Snowflake account which consumes a data share from the provider, then I think your answer makes sense. However, I'm not so sure that a reader and a consumer are considered different things.I am sorry for the spam: but here in the snowflake documentation it says a reader is still considered a consumer: https://docs.snowflake.com/en/user-guide/data-sharing-reader-create So I think your answer here is incorrect.

upvoted 1 times

**AkexS** 3 years, 3 months ago

I am sorry for the spam: but here in the snowflake documentation it says a reader is still considered a consumer: https://docs.snowflake.com/en/user-guide/data-sharing-reader-create So I think your answer here is incorrect.

upvoted 1 times

**Sid460545** 4 years, 9 months ago

Question itself is tricky. Usually Data providers provides the data to data Consumers. But a provider can also become a consumer. So can anyone please comment what should be the right answer . I think it should be False, as Data providers can share data with consumer and other providers as well

upvoted 10 times

**niho1993** Most Recent 1 year, 1 month ago

In Snowflake's Secure Data Sharing: A Data Provider can share data with multiple Consumers at the same time. Consumers can be: Other Snowflake accounts (within the same region or cross-region via replication). Reader Accounts (for consumers who don't yet have their own Snowflake account).

upvoted 1 times

**sg242** 1 year, 11 months ago

data providers can publish their data to the snowflake market.

upvoted 1 times

**hotflake** 2 years, 3 months ago

WHAT IS A SHARE? ● Named Snowflake object that encapsulates all information required to share objects within a database ● Share contains: ○ Privileges that grant access to the database, schema, or specific objects ○ Data Consumer or Reader account name ○ Shares do not contain any data

upvoted 1 times

**omiiUST** 2 years, 3 months ago

False. In Snowflake, data providers can share data not only with individual data consumers but also with multiple data consumers through various means such as: Direct Sharing: Sharing data with specific Snowflake accounts. Reader Accounts: Sharing data with organizations that do not have their own Snowflake accounts by creating and managing Reader Accounts. Snowflake Marketplace: Sharing data publicly with multiple consumers who subscribe to the data offering on the Snowflake Marketplace. Thus, data sharing in Snowflake is flexible and allows for sharing with multiple consumers in different ways.

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

B.False

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

Data Providers in Snowflake can share data with multiple Data Consumers. Snowflake's data sharing feature allows data providers to share specific databases, schemas, tables, or views with other Snowflake accounts. This sharing can be controlled at a granular level, allowing data providers to share data with multiple consumers as needed.

upvoted 1 times

**p22_nilesh** 2 years, 5 months ago

It should be A. As Reader is also a Data Consumer.

upvoted 1 times

**matias_salina** 3 years, 1 month ago

"When sharing in Snowflake, the account that shares data is called a provider, and the account that is a recipient of the data is called a consumer." "A data consumer is any account that chooses to create a database from a share made available by a data provider." "To facilitate sharing data with these consumers, you can create reader accounts. Reader accounts (formerly known as “read-only accounts”) provide a quick, easy, and cost-effective way to share data without requiring the consumer to become a Snowflake customer. Each reader account belongs to the provider account that created it" https://docs.snowflake.com/en/user-guide/data-sharing-intro

upvoted 3 times

**Dalo** 3 years, 5 months ago

False. In Snowflake, data sharing is not limited to sharing data with only a single data consumer. Snowflake allows data providers to share data with multiple data consumers or even make data publicly accessible, depending on the configured sharing settings. With Snowflake's secure data sharing capabilities, data providers can share specific datasets or entire databases with one or more data consumers. Data consumers can be given read-only access to the shared data, allowing them to query and analyze the data without being able to modify or change it. This flexible data sharing model enables collaboration and data exchange between different organizations or teams within an organization. Furthermore, Snowflake also provides granular access controls and permissions, allowing data providers to define who can access their shared data and what level of access each data consumer has. This ensures data privacy and security while enabling effective data sharing and collaboration.

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

**Selected Answer: A**

A is correct

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

A is correcthave you cleared the exam with this answers?Sorry Mayuri. But this question made my day .Nice question

upvoted 1 times

**akhandelwal680** 3 years, 5 months ago

have you cleared the exam with this answers?Sorry Mayuri. But this question made my day .Nice question

upvoted 5 times

**akhandelwal680** 3 years, 5 months ago

have you cleared the exam with this answers?

upvoted 5 times

**Rajivnb** 3 years, 3 months ago

Sorry Mayuri. But this question made my day .

upvoted 1 times

**Abhishek1180** 3 years, 4 months ago

Nice question

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

Share can be created independently

upvoted 2 times

**RameshChowdary** 3 years, 9 months ago

ANSWER IS B.

upvoted 1 times

**jurges** 3 years, 10 months ago

In the link provided in the answer you will find that by itself the data provider could share with data consumer but this data consumer could be also a data provider. The question is tricky, but by having a look at schema provided in the link answer is B

upvoted 2 times

**elvancedonzy** 4 years, 2 months ago

A Data Provider can be a Data Consumer.

upvoted 2 times

**addixion** 4 years, 7 months ago

The shared data cannot be modified so by definition it can only be consumed. The producer can change the data that is shared but in this case it is not considered shared data for the producer.

upvoted 3 times

Question #115 *Topic 1*

The fail-safe retention period is how many days?

A. 1 day

B. 7 days

C. 45 days

D. 90 days

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B.7 Days

upvoted 1 times

**Prerna5549** 3 years, 2 months ago

D is correct answer, Fail-safe is for 90 days

upvoted 2 times

**Prerna5549** 3 years, 2 months ago

D is correct answer, Fail-safe is for 90 daystime_travel up to 90 days? fail-safe 7 days?

upvoted 2 times

**nini2912** 3 years, 1 month ago

time_travel up to 90 days? fail-safe 7 days?

upvoted 2 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**seaun** 4 years ago

B is correct, Failsafe is for 7 days

upvoted 2 times

**DK69** 4 years, 5 months ago

B is correct

upvoted 2 times

Question #116 *Topic 1*

True or False: Once created, a micro-partition will never be changed.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*

A (93%)

**Siif97** Most Recent 1 year, 9 months ago

Micro-partitions are immutable

upvoted 2 times

**vhvhkkk** 1 year, 11 months ago

A. True Once a micro-partition is created in Snowflake, it cannot be modified. Micro-partitions are immutable, meaning that after data is written to them, they cannot be altered or updated directly. However, Snowflake automatically manages data storage and optimization. If data in a micro-partition needs to be updated, the system will create new micro-partitions with the updated data, and older partitions will eventually be marked for deletion during data maintenance operations like vacuuming. This immutability helps Snowflake maintain its performance and data consistency over time.

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

A.True

upvoted 1 times

**Thaiya** 3 years, 1 month ago

micro partition can be changed but not deleted.

upvoted 1 times

**ifmateus** 3 years, 2 months ago

How can we say a micro-partition can't be changed if using clustering keys we're literally changing how micro-partitions are organized? I think it's B, can't find these "immutable" references from the other answers anywhere. https://community.snowflake.com/s/article/understanding-micro-partitions-and-data-clusteringIn reclustering context, reorganizing doesn't mean that anything changes in existing micro-partitions' structure - instead the new micro-partitions are created with the same keys as in old micro-partitions before reclustering but reorganized to fit the clustering requirements. Check this article for a reference: https://docs.snowflake.com/en/user-guide/tables-clustering-keys#reclustering Check the diagram from a provided example. New micro-partitions are called "new" for a reason, also the new numbering (micro-partition 5, 6, 7, 8) should give you a clue.

upvoted 1 times

**ifmateus** 3 years, 2 months ago

Selected Answer: B

How can we say a micro-partition can't be changed if using clustering keys we're literally changing how micro-partitions are organized? I think it's B, can't find these "immutable" references from the other answers anywhere. https://community.snowflake.com/s/article/understanding-micro-partitions-and-data-clustering

upvoted 1 times

**rogerek077** 2 years, 7 months ago

In reclustering context, reorganizing doesn't mean that anything changes in existing micro-partitions' structure - instead the new micro-partitions are created with the same keys as in old micro-partitions before reclustering but reorganized to fit the clustering requirements. Check this article for a reference: https://docs.snowflake.com/en/user-guide/tables-clustering-keys#reclustering Check the diagram from a provided example. New micro-partitions are called "new" for a reason, also the new numbering (micro-partition 5, 6, 7, 8) should give you a clue.

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**jurges** 3 years, 10 months ago

The immutability of micro partition it means that there could be changes in relation to the status of the data in micro partition, so it could be updated or deleted but this is kept as information. But the data itself cannot be changed.Even if update with the clustering key it will add new micropartitions with the changes

upvoted 2 times

**kjdfsafc87456** 4 years, 2 months ago

The answer is correct as How did they do this? Let's break down how Snowflake stores data: in Snowflake, all data is stored in micro-partitions, which are contiguous units of storage that are immutable. Immutability means that once a micro-partition is written, it will never be changed

upvoted 2 times

**lhstudylh** 4 years, 8 months ago

While True, it is important to note new micro-partitions are created when new data is uploaded or data is transformed. Additional old micro-paritions are stored for time-travel support purposes

upvoted 4 times

**Sid460545** 4 years, 9 months ago

Correct, it can never be changed. They are immutable.

upvoted 1 times

Question #117 *Topic 1*

What services does Snowflake automatically provide for customers that they may have been responsible for with their on-premise system? (Choose all that apply.)

A. Installing and configuring hardware

B. Patching software

C. Physical security

D. Maintaining metadata and statistics

**Suggested Answer:** *A, B, D*

*Community vote distribution*

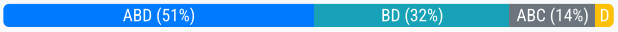


**arpit_dataguy** Highly Voted 4 years, 9 months ago

Physical security is the responsibility of the cloud vendor we have chosen. i.e. either GCP, Azure or AWS.The are not on prem though as you noted and the question refers to On-Prem. I'd argue all are correct as sometimes organizations do set up their own data centers and ensure security along with other operational aspects like electricity, water etc for cooling and maintenanceI agree its allWe are taking services from Snowflake so only snowflake will be responsible

upvoted 15 times

**arpit_dataguy** Highly Voted 4 years, 9 months ago

Selected Answer: ABD

Physical security is the responsibility of the cloud vendor we have chosen. i.e. either GCP, Azure or AWS.

upvoted 15 times

**Sonny22** 4 years, 7 months ago

The are not on prem though as you noted and the question refers to On-Prem. I'd argue all are correct as sometimes organizations do set up their own data centers and ensure security along with other operational aspects like electricity, water etc for cooling and maintenanceI agree its all

upvoted 11 times

**Sonny22** 4 years, 7 months ago

The are not on prem though as you noted and the question refers to On-Prem. I'd argue all are correct as sometimes organizations do set up their own data centers and ensure security along with other operational aspects like electricity, water etc for cooling and maintenance

upvoted 11 times

**MultiCloudIronMan** 3 years, 3 months ago

I agree its all

upvoted 1 times

**DataEngineerNeeraj1** 2 years, 1 month ago

We are taking services from Snowflake so only snowflake will be responsible

upvoted 1 times

**Santa9os** Most Recent 1 year, 2 months ago

All are correct ! “Snowflake manages all aspects of software and hardware installation and updates. This includes: • Maintaining the infrastructure (including physical and environmental security) • Managing software updates and patches • Optimizing data storage and metadata management.”

upvoted 1 times

**jeny12138** 1 year, 7 months ago

ALL TRUE

upvoted 1 times

**kentucky13caa** 1 year, 9 months ago

The correct answers are B. Patching software and D. Maintaining metadata and statistics. Here's why: B. Patching software: Snowflake handles all software patching and updates automatically, which eliminates the need for customers to manage this aspect, making it a self-managed service. D. Maintaining metadata and statistics: Snowflake automatically manages metadata and statistics, optimizing query performance without manual intervention by users. Regarding the other options: A. Installing and configuring hardware: This is handled by Snowflake, but since there is no hardware to manage (virtual or physical), it's not something customers need to do, which aligns with the self-managed nature of the platform. C. Physical security: This is the responsibility of the cloud provider (e.g., AWS, Azure, GCP), not Snowflake itself. So, the correct answers are B and D.

upvoted 1 times

**MultiCloudIronMan** 1 year, 9 months ago

Agreed

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

All options are correct

upvoted 2 times

**Yah020** 2 years, 6 months ago

This question was aslo asked in the Level Up: Snowflake Key Concepts & Architecture What common tasks for traditional on-premises database and IT staff are not required with Snowflake? Select all that apply. A Maintaining metadata B Maintaining statistics C Maintaining the physical security of a server room (key cards, door locks, etc.) D Maintaining database objects ABC correct Submit So for this question you can say all the answers are correct.

upvoted 3 times

**Heetec** 3 years ago

BD correct

upvoted 2 times

**BobFar** 3 years ago

BD is correct answer There is no hardware (virtual or physical) to select, install, configure, or manage. There is virtually no software to install, configure, or manage. Ongoing maintenance, management, upgrades, and tuning are handled by Snowflake. Physical security is the responsibility of the cloud vendor we have chosen. i.e. either GCP, Azure or AWS. https://docs.snowflake.com/en/user-guide/intro-key-concepts#

upvoted 1 times

**Marge23** 3 years, 1 month ago

BD correct ans. SNOWFLAKE is a true self-managed service, meaning: • There is no hardware (virtual or physical) to select, install, configure, or manage. • There is virtually no software to install, configure, or manage. • Ongoing maintenance, management, upgrades, and tuning are handled by Snowflake.

upvoted 3 times

**AbhaySAP** 3 years, 4 months ago

BD - As Snowflake is SAAS based, and Hardware and Physical Security doesn't come under SAAS Snowflake is a true self-managed service, meaning: There is no hardware (virtual or physical) to select, install, configure, or manage. There is virtually no software to install, configure, or manage. Ongoing maintenance, management, upgrades, and tuning are handled by Snowflake.When choosing cluster size, SF will handle the hardware provisioning. SaaS is provided for end users, although it is SF's responsibility to take care of hardware, software, and metadata.

upvoted 2 times

**AbhaySAP** 3 years, 4 months ago

BD - As Snowflake is SAAS based, and Hardware and Physical Security doesn't come under SAAS Snowflake is a true self-managed service, meaning: There is no hardware (virtual or physical) to select, install, configure, or manage. There is virtually no software to install, configure, or manage. Ongoing maintenance, management, upgrades, and tuning are handled by Snowflake.

upvoted 2 times

**PrashantGupta1616** 2 years, 4 months ago

When choosing cluster size, SF will handle the hardware provisioning. SaaS is provided for end users, although it is SF's responsibility to take care of hardware, software, and metadata.

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

ABD are correct

upvoted 2 times

**bibisa** 3 years, 7 months ago

B. Patching software D. Maintaining metadata and statistics Snowflake is a cloud-based data warehousing platform that eliminates the need for customers to manage hardware infrastructure, physical security, and related tasks. These responsibilities are handled by Snowflake's cloud infrastructure provider, and Snowflake also automatically manages the patching of software and maintenance of metadata and statistics for its customers. Option A (Installing and configuring hardware) and Option C (Physical security) are not provided by Snowflake since they are handled by the cloud infrastructure provider.

upvoted 4 times

**SD5713** 3 years, 7 months ago

So can someone help me which one is correct?

upvoted 1 times

**KarBiswa** 3 years, 7 months ago

B and D are the correct options https://docs.snowflake.com/en/user-guide/intro-key-concepts#:~:text=There%20is%20no,handled%20by%20Snowflake.

upvoted 1 times

**EmiB** 3 years, 8 months ago

All 4 Snowflake "provide for customers" - doesn't matter if SF itself or Cloud provider. Customer has contract with Snowflake, not with the Cloud provider.

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

No physical security as it is cloud based and no hardware usage

upvoted 3 times

Question #118 *Topic 1*

Which of the following statements would be used to export/unload data from Snowflake?

A. COPY INTO @stage

B. EXPORT TO @stage

C. INSERT INTO @stage

D. EXPORT_TO_STAGE(stage => @stage, select => 'select * from t1');

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Altria** Most Recent 1 year, 1 month ago

COPY INTO command is used for unloading.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A Copy into @stage

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 2 times

**Mayuri917233** 3 years, 5 months ago

A is correct

upvoted 1 times

**lhstudylh** 4 years, 8 months ago

Correct answer

upvoted 1 times

Question #119 *Topic 1*

True or False: A 4X-Large Warehouse may, at times, take longer to provision than a X-Small Warehouse.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*




**Dilbert** Highly Voted 3 years, 8 months ago

Selected Answer: A

Signed up for free Snowflake account. Created X-Small and 4X Virtual Warehouse using web interface. X-Small took 96 ms, while 4X took 105 ms. You can validate it in Activity -> Query History.

upvoted 7 times

**Dilbert** Highly Voted 3 years, 8 months ago

Signed up for free Snowflake account. Created X-Small and 4X Virtual Warehouse using web interface. X-Small took 96 ms, while 4X took 105 ms. You can validate it in Activity -> Query History.I tried, with X-small & 6X-Large. Interestingly X-small took 4 ms more than 6X-Large

upvoted 10 times

**omiiUST** Most Recent 2 years, 3 months ago

I tried, with X-small & 6X-Large. Interestingly X-small took 4 ms more than 6X-Large

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

Due to large number of compute resourecs it take longer time,A.True

upvoted 2 times

**Mike722** 2 years, 8 months ago

https://docs.snowflake.com/en/user-guide/warehouses-considerations. You require the warehouse to be available with no delay or lag time. Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer.

upvoted 4 times

**sandy91** 2 years, 9 months ago

Answer is A. You require the warehouse to be available with no delay or lag time. Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer.

upvoted 2 times

**BobFar** 3 years ago

Answer is A. You require the warehouse to be available with no delay or lag time. Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer. https://docs.snowflake.com/en/user-guide/warehouses-considerations

upvoted 1 times

**Marge23** 3 years, 1 month ago

Correct answer is B. no reference that you can find that 4x is provisioned longer than xs w.

upvoted 1 times

**Marge23** 3 years, 1 month ago

Correct answer is B. no reference that you can find that 4x is provisioned longer than xs w.No, correct Answer is A. Here is the reference: You require the warehouse to be available with no delay or lag time. Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer. https://docs.snowflake.com/en/user-guide/warehouses-considerations

upvoted 1 times

**BobFar** 3 years ago

No, correct Answer is A. Here is the reference: You require the warehouse to be available with no delay or lag time. Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer. https://docs.snowflake.com/en/user-guide/warehouses-considerations

upvoted 3 times

**Mayuri917233** 3 years, 5 months ago

B is correct

upvoted 1 times

**learner2023** 3 years, 6 months ago

You require the warehouse to be available with no delay or lag time. Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer.

upvoted 3 times

**sdnv89** 3 years, 11 months ago

Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer.So the answer is A and not B... depending on the size of the warehouse ... => B

upvoted 2 times

**sdnv89** 3 years, 11 months ago

Selected Answer: B

Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer.

upvoted 2 times

**YokoSumiGaeshi** 3 years, 8 months ago

So the answer is A and not B

upvoted 7 times

**EmiB** 3 years, 8 months ago

... depending on the size of the warehouse ... => B

upvoted 3 times

**greenlever** 4 years ago

Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer.

upvoted 4 times

**distopic** 4 years, 4 months ago

Selected Answer: B

Only 5x and 6x can need more time because are in preview

upvoted 2 times

**distopic** 4 years, 4 months ago

Only 5x and 6x can need more time because are in previewNo the answer is A. You require the warehouse to be available with no delay or lag time. Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer.

upvoted 2 times

**BobFar** 3 years ago

No the answer is A. You require the warehouse to be available with no delay or lag time. Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer.

upvoted 1 times

**S_J** 4 years, 5 months ago

False Provisioning the larger warehouse sizes 5X-Large and 6X-Large may take slightly longer while in preview. https://docs.snowflake.com/en/user-guide/warehouses-overview.html

upvoted 2 times

**Sid460545** 4 years, 9 months ago

It is true, it can take longer time to provision

upvoted 4 times

**arpit_dataguy** 4 years, 9 months ago

The answer is true. You can experiment the same with snowflake UI.

upvoted 2 times

Question #120 *Topic 1*

How would you determine the size of the virtual warehouse used for a task?

A. Root task may be executed concurrently (i.e. multiple instances), it is recommended to leave some margins in the execution window to avoid missing instances of execution

B. Querying (SELECT) the size of the stream content would help determine the warehouse size. For example, if querying large stream content, use a larger warehouse size

C. If using the stored procedure to execute multiple SQL statements, it's best to test run the stored procedure separately to size the compute resource first

D. Since task infrastructure is based on running the task body on schedule, it's recommended to configure the virtual warehouse for automatic concurrency handling using Multi-cluster warehouse (MCW) to match the task schedule

**Suggested Answer:** *C*

*Community vote distribution*

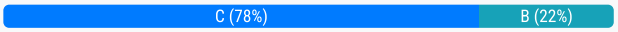


**SV1122** Highly Voted 3 years, 10 months ago

"Multi-cluster warehouses are best utilized for scaling resources to improve concurrency for users/queries. They are not as beneficial for improving the performance of slow-running queries or data loading. For these types of operations, resizing the warehouse provides more benefits." Snowflake representative explained this here https://community.snowflake.com/s/question/0D50Z00009F8nPTSAZ/how-to-choose-the-right-virtual-warehouse-size-in-snowflake-for-your-workloadThe article was really self-explanatory. Thank you.

upvoted 5 times

**SilviaS** 10 months, 1 week ago

The article was really self-explanatory. Thank you.

upvoted 1 times

**SV1122** 3 years, 10 months ago

Selected Answer: C

"Multi-cluster warehouses are best utilized for scaling resources to improve concurrency for users/queries. They are not as beneficial for improving the performance of slow-running queries or data loading. For these types of operations, resizing the warehouse provides more benefits."

Snowflake representative explained this here https://community.snowflake.com/s/question/0D50Z00009F8nPTSAZ/how-to-choose-the-right-virtual-warehouse-size-in-snowflake-for-your-workload

upvoted 4 times

**alexhparreira** Most Recent 5 months ago

Why Option B is correct: Tasks are frequently used to process change data capture (CDC) pipelines fed by Snowflake Streams. When sizing a warehouse for a task, your main performance metric is the volume of data waiting to be processed. By checking the stream size (e.g., using SYSTEM$STREAM_HAS_DATA or checking row counts), you can dynamically gauge or establish the baseline warehouse size required to ingest that specific workload efficiently.

upvoted 1 times

**niho1993** 1 year, 1 month ago

B. Querying (SELECT) the size of the stream content would help determine the warehouse size. □ Correct. If your task is processing large streams (big batch of changes), you need more compute → bigger warehouse. If stream is small, XS/S is fine. C. If using the stored procedure to execute multiple SQL statements, it's best to test run the stored procedure separately to size the compute resource first. □ Correct. Since stored procs can have multiple steps (DML, transformations), you should benchmark them manually to pick the right warehouse size.

upvoted 1 times

**mandar_007** 2 years ago

B For queries in large-scale production environments, larger warehouse sizes (Large, X-Large, 2X-Large, etc.) may be more cost effective. https://docs.snowflake.com/en/user-guide/warehouses-considerations

upvoted 1 times

**Mayuri917233** 3 years, 5 months ago

C is correct

upvoted 1 times

**grabc** 4 years, 9 months ago

C should be correct

upvoted 1 times

**arpit_dataguy** 4 years, 9 months ago

C is the right answer as D talks about concurrency.

upvoted 1 times

**corentin** 4 years, 9 months ago

I think B

upvoted 4 times

Question #121 *Topic 1*

The Information Schema and Account Usage Share provide storage information for which of the following objects? (Choose three.)

A. Users

B. Tables

C. Databases

D. Internal Stages

**Suggested Answer:** *B, C, D*

*Community vote distribution*

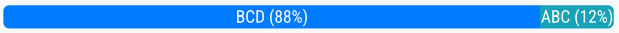


**William_20** Highly Voted 4 years, 2 months ago

Selected Answer: BCD

Couldn't find any information where storage of a user is explained/showed. Tables in Account Usage view related to storage are (database, stage, table): DATABASE_STORAGE_USAGE_HISTORY, STAGE_STORAGE_USAGE_HISTORY, STORAGE_USAGE, TABLE_STORAGE_METRICS

upvoted 12 times

**William_20** Highly Voted 4 years, 2 months ago

Couldn't find any information where storage of a user is explained/showed. Tables in Account Usage view related to storage are (database, stage, table): DATABASE_STORAGE_USAGE_HISTORY, STAGE_STORAGE_USAGE_HISTORY, STORAGE_USAGE, TABLE_STORAGE_METRICSWe can get User info too. The Account Usage view can be used to query a list of all users in the account. The data is retained for 365 days (1 year). https://docs.snowflake.com/en/sql-reference/account-usage/users.htmlProbably BCD Because users are not "storage information" as mentioned in the question

upvoted 12 times

**SV1122** 3 years, 10 months ago

We can get User info too. The Account Usage view can be used to query a list of all users in the account. The data is retained for 365 days (1 year). https://docs.snowflake.com/en/sql-reference/account-usage/users.html

upvoted 3 times

**Def21** 3 years, 4 months ago

Probably BCD Because users are not "storage information" as mentioned in the question

upvoted 3 times

**visionary4ever** Most Recent 1 year, 2 months ago

Users don’t “consume storage”; they consume compute when they run queries.

upvoted 1 times

**Santa9os** 1 year, 2 months ago

Users are security principals, not storage objects. Storage usage views do not report storage per user.

upvoted 1 times

**37ceea2** 1 year, 9 months ago

It is about storage information. We can see information related to users, such as queries accessed by a particular user, or how many logins and when did the user logged, but there is not such case of storage information related to an used.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

BCD ARE CORRECT

upvoted 1 times

**_yyukta** 2 years, 8 months ago

BCD are correct

upvoted 1 times

**snair25** 2 years, 12 months ago

Internal Stages comsume storage and the storage consumed can be found in the SNOWFLAKE.ACCOUNT_USAGE.STAGE_STORAGE_USAGE_HISTORY. Storage consumed by tables and databases can be found in SNOWFLAKE.ACCOUNT_USAGE.TABLE_STORAGE_METRICS in which we have the table name and the database name included in them.

upvoted 2 times

**Amitsnowflake** 3 years, 1 month ago

It store view for all object contained in the database as well as view for account level objects such as role warehouse It store table functions for historical and usage data across your account. Select * from table (infirmation_schema.login_history_by_user(USER_NAME => 'user')) gives you login user information

upvoted 1 times

**jAtlas7** 3 years, 1 month ago

**Selected Answer: BCD**

The question is related to "Storage Information". Users do not have storage either as an object or data they hold (roles own items). Therefore users is not an option.

upvoted 2 times

**Amitsnowflake** 3 years, 1 month ago

A is also one of the correct option.There are table functions stored in information schema which can be used to get historical information for storage, warehouse, user logins and queries. Login_history, login_history_by_user (data retention 7 days)

upvoted 1 times

**Marge23** 3 years, 1 month ago

ABC SNOWFLAKE VIEWS: USERS, DATABASES, SCHEMATA AND TABLES NO INTERNAL STAGES.

upvoted 1 times

**jAtlas7** 3 years, 1 month ago

The question is related to "Storage Information". Users do not have storage either as an object or data they hold (roles own items). Therefore users is not an option.makes sense

upvoted 2 times

**VLMBL_94** 2 years, 2 months ago

makes sense

upvoted 1 times

**Marge23** 3 years, 1 month ago

Correct answer is ABC. users is the one holding for the account.

upvoted 1 times

**Hopefan** 3 years, 5 months ago

Snowflake Information Schema: Stages, External Tables and Pipes,Databases, Schema, Tables and Views,Columns and Constraints and many more

upvoted 1 times

**Hopefan** 3 years, 5 months ago

I belive all fours, as I saw in the documentation. https://docs.snowflake.com/en/sql-reference/account-usage

upvoted 3 times

**SteveCrown** 3 years, 6 months ago

Well, having painstakingly checked the 2 schema, both information_schema and account_usage hvae Tables, Databases and stages. Though there is the USERS VIEW in account_usage schema, there is no such view in INFORMATION _SCHEMA. Therefore, I will argue that B,C,D are correct.

upvoted 2 times

**AnkitaS26** 3 years, 10 months ago

BCD. Could not find user information in Information Schema.

upvoted 3 times

**manish6007** 3 years, 6 months ago

Ans is ABC. As I can see only external stages while running below query: select * from SNOWFLAKE.ACCOUNT_USAGE.STAGES; I am able to get the user details using below query select * from SNOWFLAKE.ACCOUNT_USAGE.USERS ;

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

No user level

upvoted 2 times

**NEMMELA** 3 years, 2 months ago

correct

upvoted 1 times

**SV1122** 3 years, 10 months ago

I found information that supports all answers ABCD. This is confusing.

upvoted 1 times

**BigDataBB** 3 years, 9 months ago

select * from snowflake.information_schema.views where table_name like '%STORAGE%' and table_schema in ('ACCOUNT_USAGE','INFORMATION_SCHEMA');

TABLE_CATALOG TABLE_SCHEMA TABLE_NAME SNOWFLAKE ACCOUNT_USAGE DATABASE_STORAGE_USAGE_HISTORY SNOWFLAKE ACCOUNT_USAGE STAGE_STORAGE_USAGE_HISTORY SNOWFLAKE ACCOUNT_USAGE STORAGE_USAGE SNOWFLAKE ACCOUNT_USAGE TABLE_STORAGE_METRICS SNOWFLAKE INFORMATION_SCHEMA TABLE_STORAGE_METRICS

upvoted 2 times

Question #122 *Topic 1*

What is the default File Format used in the COPY command if one is not specified?

A. CSV

B. JSON

C. Parquet

D. XML

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A. CSV

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. CSV

upvoted 1 times

**Marge23** 3 years, 1 month ago

csv correct answer.

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**fahfouhi94** 3 years, 8 months ago

csv is default format

upvoted 1 times

**allthingsdata** 4 years, 4 months ago

A is the correct answer

upvoted 1 times

Question #123 *Topic 1*

True or False: Reader Accounts are able to extract data from shared data objects for use outside of Snowflake.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*




**Santa9os** Most Recent 1 year, 2 months ago

"Reader accounts can be used to query and view shared data, but cannot extract data out of Snowflake (e.g. unload data into files, or copy data into another database)."

upvoted 1 times

**jeny12138** 1 year, 7 months ago

Reader Accounts are special Snowflake accounts created and managed by a data provider to let consumers query shared data without needing their own Snowflake account. These accounts are intended for read-only access to shared data within Snowflake. They cannot export or unload data to external locations like S3, Azure Blob Storage, or even download locally.

upvoted 2 times

**PrasantSadangi** 1 year, 7 months ago

You cannot perform the following tasks in a reader account: - Set a data metric function on objects in the reader account. - Upload new data. - Modify existing data. - Unload data using a storage integration. "However, you can use the COPY INTO command with your connection credentials to unload data into a cloud storage location."

upvoted 2 times

**MultiCloudIronMan** 1 year, 9 months ago

I think its false, copy into a doesnt mean you can use it because it can be a metadata copy.

upvoted 1 times

**56d43d6** 1 year, 10 months ago

A: https://docs.snowflake.com/en/user-guide/data-sharing-reader-create#what-is-restricted-allowed-in-a-reader-account

upvoted 2 times

**capt2101akash** 2 years, 4 months ago

Reader accounts can unload data to storage integration using COPY INTO

upvoted 1 times

**pranalig** 2 years, 4 months ago

A. True

upvoted 1 times

**bot314** 2 years, 4 months ago

you cannot execute the following commands in a reader account: INSERT UPDATE DELETE MERGE COPY INTO CREATE MASKING POLICY CREATE PIPE CREATE ROW ACCESS POLICY CREATE SHARE CREATE STAGE SHOW PROCEDURES

upvoted 2 times

**PrashantGupta1616** 2 years, 5 months ago

A. True Reader Accounts in Snowflake have the ability to extract data from shared data objects for use outside of Snowflake. This capability allows organizations to securely share data with external parties while still maintaining control over access and usage.

upvoted 1 times

**Alex_Ovidiu** 2 years, 7 months ago

False: Reader Accounts are not able to extract data from shared data objects for use outside of Snowflake. Here's why: Snowflake's Secure Data Sharing with Reader Accounts provides a mechanism to grant controlled access to data without requiring the recipient to have a full Snowflake account. Reader Accounts can query and analyze the shared data objects (tables, views, etc.) within the Snowflake environment. However, they cannot directly extract the data and download it for use outside of Snowflake. This restriction helps ensure data security and governance. Data providers maintain control over who can access the data and how it's used. While users in reader accounts can gain insights from the data, they cannot export it for potential misuse or unauthorized sharing.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A True

upvoted 1 times

**lol632** 2 years, 11 months ago

What is restricted/allowed in a reader account? A reader account is intended primarily for querying data shared by the provider of the account. You can work with data, for example, by creating materialized views. You cannot perform the following tasks in a reader account: - Uploading new data. - Modifying existing data. - Unloading data using a storage integration. However, you can use the COPY INTO command with your connection credentials to unload data into a cloud storage location. So the answer is A, you can export data into a cloud storage location. Link: https://docs.snowflake.com/en/user-guide/data-sharing-reader-create

upvoted 4 times

**Amitsnowflake** 3 years ago

yes, Reader account can download data

upvoted 1 times

**Gwo** 3 years, 1 month ago

A. Reader accounts are Snowflake accounts that are created and managed by data providers on behalf of their consumers. They allow data consumers to: Access and query data shared by the provider of the account Explore and export their data Create customized exports and materialized views

upvoted 2 times

**Roohee** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/data-sharing-reader-create Reader accounts (formerly known as “read-only accounts”) enable providers to share data with consumers who are not already Snowflake customers, without requiring the consumers to become Snowflake customers.

upvoted 1 times

**surya_lolla** 3 years, 3 months ago

Unloading data using a storage integration. However, you can use the COPY INTO command with your connection credentials to unload data into a cloud storage location. https://docs.snowflake.com/en/user-guide/data-sharing-reader-create

upvoted 1 times

**MadhaviS** 3 years, 2 months ago

https://community.snowflake.com/s/question/0D5Do00000H2a5fKAB/hello-expertsis-it-possible-for-reader-accounts-to-extract-data-from-shared-data-objects-for-outside-use-of-snowflakemany-thanks-in-advance

upvoted 1 times

Question #124 *Topic 1*

True or False: You can define multiple columns within a clustering key on a table.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*

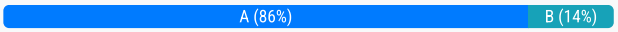


**Mallikharjuna452** Most Recent 2 years, 3 months ago

Yes it is true

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A True

upvoted 1 times

**SV1122** 3 years, 10 months ago

-- cluster by base columns create or replace table t1 (c1 date, c2 string, c3 number) cluster by (c1, c2); https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html

upvoted 1 times

**nicholascz** 4 years ago

"A single clustering key can contain one or more columns or expressions." https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html

upvoted 2 times

**William_20** 4 years, 2 months ago

https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html#strategies-for-selecting-clustering-keys. "A single clustering key can contain one or more columns or expressions."

upvoted 2 times

**kjdfsafc87456** 4 years, 2 months ago

WE CANNOT DO THIS AS THE FUNCTION IS DEPRECIATED

upvoted 1 times

**rkarthik0789** 4 years, 4 months ago

Can define 3 or 4 combination of columns, mostly using the columns used in Join and filter

upvoted 1 times

Question #125 *Topic 1*

True or False: Snowflake enforces unique, primary key, and foreign key constraints during DML operations.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**BigDataBB** Highly Voted 3 years, 9 months ago

Snowflake supports defining and maintaining constraints, but does not enforce them, except for NOT NULL constraints, which are always enforced.

upvoted 5 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B.False

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B false

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**sdnv89** 3 years, 11 months ago

Only NOT NULL is enforced. https://docs.snowflake.com/en/sql-reference/constraints.html

upvoted 2 times

**Sid460545** 4 years, 9 months ago

It doesn't enforce.. But Not null is

upvoted 3 times

Question #126 *Topic 1*

True or False: Loading data into Snowflake requires that source data files be no larger than 16MB.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**jeny12138** Most Recent 1 year, 7 months ago

The number of load operations that run in parallel can't exceed the number of data files to be loaded. To optimize the number of parallel operations for a load, we recommend aiming to produce data files roughly 100-250 MB (or larger) in size compressed.

upvoted 2 times

**DataEngineerNeeraj1** 2 years, 1 month ago

source data files can be 100-250 MB compressed, 16 MB mentioned is size of each rows after getting loaded into table

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

B.False

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B False

upvoted 1 times

**aqmdla2002** 4 years, 3 months ago

https://docs.snowflake.com/en/user-guide/data-load-considerations-prepare.html

upvoted 1 times

**Sid460545** 4 years, 9 months ago

16 Mb is default. It can support till 5GBAlso 16 MB is considered "best practice"

upvoted 3 times

**Sid460545** 4 years, 9 months ago

16 Mb is default. It can support till 5GB

upvoted 3 times

**lhstudylh** 4 years, 8 months ago

Also 16 MB is considered "best practice"

upvoted 2 times

**arpit_dataguy** 4 years, 9 months ago

By default, COPY INTO location statements separate table data into a set of output files to take advantage of parallel operations. The maximum size for each file is set using the MAX_FILE_SIZE copy option. The default value is 16777216 (16 MB) but can be increased to accommodate larger files. The maximum file size supported is 5 GB for Amazon S3, Google Cloud Storage, or Microsoft Azure stages. To unload data to a single output file (at the potential cost of decreased performance), specify the SINGLE = true copy option in your statement. You can optionally specify a name for the file in the path.

upvoted 3 times

Question #127 *Topic 1*

True or False: A Virtual Warehouse can be resized while suspended.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*

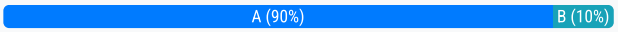


**56d43d6** Most Recent 1 year, 10 months ago

https://docs.snowflake.com/en/user-guide/warehouses-considerations#warehouse-resizing-improves-performance

upvoted 2 times

**Mochin** 1 year, 10 months ago

It is False as you cannot resize a virtual warehouse while it is suspended, you need to have a active warehouse to provide the compute clustersnot accurate. https://docs.snowflake.com/en/user-guide/warehouses-considerations#warehouse-resizing-improves-performance a wh can be resized at any time.

upvoted 1 times

**Mochin** 1 year, 10 months ago

Selected Answer: B

It is False as you cannot resize a virtual warehouse while it is suspended, you need to have a active warehouse to provide the compute clusters

upvoted 1 times

**56d43d6** 1 year, 10 months ago

not accurate. https://docs.snowflake.com/en/user-guide/warehouses-considerations#warehouse-resizing-improves-performance a wh can be resized at any time.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A.True

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. True

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**chizo** 3 years, 7 months ago

Ans is true

upvoted 2 times

**SV1122** 3 years, 10 months ago

Yes, it can be resized both when it is suspended or active (even when queries are running).

upvoted 3 times

**datamech001** 4 years ago

Answer A

upvoted 1 times

Question #128 *Topic 1*

True or False: When you create a custom role, it is a best practice to immediately grant that role to ACCOUNTADMIN.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

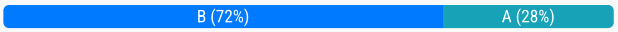


**addixion** Highly Voted 4 years, 7 months ago

100% false. SYSADMIN should immediately be granted the role, then ACCOUNTADMIN will automatically have the grant through SYSADMIN.

upvoted 16 times

**mandar_007** Most Recent 2 years ago

B Instead, we recommend creating a hierarchy of roles aligned with business functions in your organization and ultimately assigning these roles to the SYSADMIN role.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B.False

upvoted 1 times

**rogerek077** 2 years, 7 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-overview#roles "When creating roles that will serve as the owners of securable objects in the system, Snowflake recommends creating a hierarchy of custom roles, with the top-most custom role assigned to the system role SYSADMIN." B is correct according to that article

upvoted 4 times

**GRKaushal** 2 years, 8 months ago

We don't have to granted the role to Accountadmin Its automatically have throught the SysADMIN role

upvoted 2 times

**_yyukta** 2 years, 8 months ago

B. False

upvoted 1 times

**panksoo** 3 years, 7 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-considerations Important: By default, not even the ACCOUNTADMIN role can modify or drop objects created by a custom role. The custom role must be granted to the ACCOUNTADMIN role directly or, preferably, to another role in a hierarchy with the SYSADMIN role as the parent. The SYSADMIN role is managed by the ACCOUNTADMIN role.correct A

upvoted 3 times

**panksoo** 3 years, 7 months ago

Selected Answer: A

https://docs.snowflake.com/en/user-guide/security-access-control-considerations Important: By default, not even the ACCOUNTADMIN role can modify or drop objects created by a custom role. The custom role must be granted to the ACCOUNTADMIN role directly or, preferably, to another role in a hierarchy with the SYSADMIN role as the parent. The SYSADMIN role is managed by the ACCOUNTADMIN role.

upvoted 3 times

**NEMMELA** 3 years, 2 months ago

correct A

upvoted 1 times

**ShagunMittal** 3 years, 7 months ago

keyword in the question is "immediately" and that makes the correct answer as False. Granting custom role to ACCOUNTADMIN is not the only option. It can be assigned to role hierarchy with SYSADMIN as parent also

upvoted 3 times

**chizo** 3 years, 7 months ago

True ans A Important By default, not even the ACCOUNTADMIN role can modify or drop objects created by a custom role. The custom role must be granted to the ACCOUNTADMIN role directly or, preferably, to another role in a hierarchy with the SYSADMIN role as the parent. The SYSADMIN role is managed by the ACCOUNTADMIN role.

upvoted 1 times

**ShagunMittal** 3 years, 7 months ago

Best practice is to assign it to ACCOUNTADMIN;

upvoted 1 times

**c0d3g** 3 years, 7 months ago

By default, not even the ACCOUNTADMIN role can modify or drop objects created by a custom role. The custom role must be granted to the ACCOUNTADMIN role directly or, preferably, to another role in a hierarchy with the SYSADMIN role as the parent. The SYSADMIN role is managed by the ACCOUNTADMIN role. https://docs.snowflake.com/en/user-guide/security-access-control-considerations#managing-custom-roles Even it is mentioned preferably SYSADMIN, the idea is to grant custom role to ACCOUNTADMIN.

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

The custom role must be granted to the ACCOUNTADMIN role directly or, preferably, to another role in a hierarchy with the SYSADMIN role as the parent.

upvoted 2 times

**William_20** 4 years, 2 months ago

Selected Answer: A

According to the docs: Accountadmin. https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html#managing-custom-roles

upvoted 1 times

**William_20** 4 years, 2 months ago

According to the docs: Accountadmin. https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html#managing-custom-rolesSorry, should be B: False. It states "preferably, to another role in a hierarchy with the SYSADMIN role as the parent.". Can't remove my previous comment and vote ...

upvoted 1 times

**William_20** Highly Voted 4 years, 2 months ago

Sorry, should be B: False. It states "preferably, to another role in a hierarchy with the SYSADMIN role as the parent.". Can't remove my previous comment and vote ...

upvoted 5 times

**Casual** 4 years, 8 months ago

B is correct

upvoted 2 times

**grabc** 4 years, 8 months ago

FALSE, Should be assigned to SYSADMIN role

upvoted 2 times

**grabc** 4 years, 9 months ago

The best practice is to assign it to SYSADMIN role, B is correct

upvoted 1 times

**arpit_dataguy** 4 years, 9 months ago

It is TRUE.

upvoted 1 times

Question #129 *Topic 1*

Which of the following accurately represents how a table fits into Snowflake's logical container hierarchy?

A. Account -> Schema -> Database -> Table

B. Account -> Database -> Schema -> Table

C. Database -> Schema -> Table -> Account

D. Database -> Table -> Schema -> Account

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B represents correct order

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Account -> Database -> Schema -> Table

upvoted 1 times

**Marge23** 3 years, 1 month ago

B is correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**SV1122** 3 years, 10 months ago

Account -> Database -> Schema -> Table

upvoted 1 times

**learnazureportal** 4 years, 6 months ago

Answer is correct

upvoted 2 times

Question #130 *Topic 1*

True or False: All Snowflake table types include fail-safe storage.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**EmiB** Highly Voted 3 years, 7 months ago

B. False (only Permanent table has Fail-safe functionality)

upvoted 5 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B.False(Only permanent table)

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: B

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. False

upvoted 1 times

**aqmdla2002** 4 years, 3 months ago

Selected Answer: B

Transient and temporary not

upvoted 4 times

**aqmdla2002** 4 years, 3 months ago

Transient and temporary notExternal tables also do not have Fail-Safe (just saying)

upvoted 4 times

**OTE** 3 years, 7 months ago

External tables also do not have Fail-Safe (just saying)

upvoted 3 times

**lhstudylh** 4 years, 8 months ago

Correct answer Temporary tables do not have fail-safe functionalityTransient tables as well

upvoted 2 times

**lhstudylh** 4 years, 8 months ago

Selected Answer: B

Correct answer

Temporary tables do not have fail-safe functionality

upvoted 2 times

**addixion** 4 years, 7 months ago

Transient tables as well

upvoted 2 times

Question #131 *Topic 1*

What are two ways to create and manage Data Shares in Snowflake? (Choose two.)

A. Via the Snowflake Web Interface (UI)

B. Via the DATA_SHARE=TRUE parameter

C. Via SQL commands

D. Via Virtual Warehouses

**Suggested Answer:** *A, C*

*Community vote distribution*

AC (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A and C

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: AC

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A and C are correct

upvoted 1 times

**Marge23** 3 years, 1 month ago

AC is the correct answer

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**arpit_dataguy** 4 years, 9 months ago

Yes, A & C is the correct answer.

upvoted 3 times

Question #132 *Topic 1*

True or False: Fail-safe can be disabled within a Snowflake account.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*

B (93%)

**JustVicky** Highly Voted 4 years, 9 months ago

False : Separate and distinct from Time Travel, Fail-safe ensures historical data is protected in the event of a system failure or other catastrophic event, e.g. a hardware failure or security breach. Fail safe feature cannot be enabled or disabled from the user end .

upvoted 6 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B.False(it is built in featur

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: AC

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B false

upvoted 1 times

**Marge23** 3 years, 1 month ago

True. you can disabled by setting your parameter to 0.That is time travel that you can set to 0

upvoted 1 times

**Marge23** 3 years, 1 month ago

True. you can disabled by setting your parameter to 0.

upvoted 1 times

**matias_salina** Highly Voted 3 years, 1 month ago

That is time travel that you can set to 0

upvoted 5 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**ANNIEKOSHY** 3 years, 4 months ago

The account administrators in Snowflake have the ability to disable fail-safe storage if needed. Disabling fail-safe storage can be done at the account level, and it affects all tables within the account. This action should be taken with caution as it removes the ability to recover data to any point in time within the retention period.

upvoted 1 times

**RajuNaik** 3 years, 6 months ago

we can not disable fail safe

upvoted 1 times

**EmiB** 3 years, 7 months ago

B. False (you can't configure it)

upvoted 1 times

**KORIR2_VIOLA** 3 years, 8 months ago

THIS IS NOT CORRECT FAIL SAVE IS NOT CONFIGARABLE

upvoted 2 times

**KarBiswa** 3 years, 8 months ago

default

upvoted 2 times

**ajay_1233456** 3 years, 9 months ago

B is correct answer, as fail-safe is non configurabl

upvoted 1 times

**baig123** 3 years, 10 months ago

B is correct answer, as fail-safe is non configurable, however, we can say indirectly yes because transient database/schema and tables are not having fail-safe option.

upvoted 1 times

**BungyTex** 4 years, 4 months ago

Failsafe cannot be disabled

upvoted 3 times

**Surbin** 4 years, 7 months ago

Fail safe can't be disabled

upvoted 3 times

**addixion** 4 years, 7 months ago

You cannot configure fail-safe. You can use transient database if you don't want to use it.

upvoted 4 times

**Sid460545** 4 years, 9 months ago

False, you cannot modify fail safe

upvoted 1 times

Question #133 *Topic 1*

True or False: It is possible for a user to run a query against the query result cache without requiring an active Warehouse.

A. True

B. False

**Suggested Answer:** *A*

*Community vote distribution*




**arpit_dataguy** Highly Voted 4 years, 9 months ago

Query result cache is all about fetching the data from cloud services layer and saving the cost by not running the virtual warehouse.

upvoted 8 times

**bor4un** Most Recent 1 year, 9 months ago

Correct

upvoted 1 times

**ABRocks** 2 years, 2 months ago

It's True; you do not need an ACTIVE warehouse to execute the query result cache.

upvoted 1 times

**VLMBL_94** 2 years, 2 months ago

User cant point the query to run using the cache.

upvoted 2 times

**vivekmani2021** 2 years, 3 months ago

the query cache is dropped when the warehouse is suspended , so you need an active warehouse

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A.True

upvoted 1 times

**pranalig** 2 years, 4 months ago

False is correct without selection of warehouse an query cannot be executed

upvoted 2 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: A

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. True

upvoted 1 times

**Jules27** 2 years, 9 months ago

The answer is B - False - https://docs.snowflake.com/en/user-guide/performance-query-warehouse-cache?utm_source=snowscope&utm_medium=serp&utm_term=query+cache The auto-suspension setting of the warehouse can have a direct impact on query performance because the cache is dropped when the warehouse is suspended.Your response doesn't actually answer the question. Yes if the wh has been suspended and the cache is dropped, sure, the cache is no longer available for re-use. but the answer to the question is A: True.

upvoted 2 times

**Jules27** 2 years, 9 months ago

The answer is B - False - https://docs.snowflake.com/en/user-guide/performance-query-warehouse-cache?utm_source=snowscope&utm_medium=serp&utm_term=query+cache The auto-suspension setting of the warehouse can have a direct impact on query performance because the cache is dropped when the warehouse is suspended.

upvoted 2 times

**56d43d6** 1 year, 10 months ago

Your response doesn't actually answer the question. Yes if the wh has been suspended and the cache is dropped, sure, the cache is no longer available for re-use. but the answer to the question is A: True.

upvoted 1 times

**Rajivnb** 3 years, 3 months ago

The answer is B. Question is not about can a user query utilize the data from cache. Its about running query against a stored query result cahce. That is not possible. User has no option to do anything with cached results.

upvoted 1 times

**Leinho** 3 years, 5 months ago

You cannot run a query against cache, if you run the exact same query that some else previously ran, snowflake has the capability to retrieve the result from cache without running a warehouse, but non user can see others users query results.

upvoted 1 times

**SravantiJana** 3 years, 5 months ago

A query can use the query result cache but can not query from it.

upvoted 1 times

**fahfouhi94** 3 years, 8 months ago

the correct answer is A , daba can be no changed last 24h so possible

upvoted 1 times

**systematics78** 3 years, 9 months ago

A is not true. The SELECT statement requires warehouse. The result set can be retrieved if data is not changed, not use some functions and etc.. but in any case, you can't use a select statement without a warehouseI upvoted this because it intuitively makes sense. But then I tested it in Snowflake and I was able to retrieve a result from the result cache using a SELECT while using a suspended warehouse (without auto resume). So A is indeed the correct answer.

upvoted 2 times

**systematics78** 3 years, 9 months ago

A is not true. The SELECT statement requires warehouse. The result set can be retrieved if data is not changed, not use some functions and etc.. but in any case, you can't use a select statement without a warehouse

upvoted 2 times

**YokoSumiGaeshi** 3 years, 8 months ago

I upvoted this because it intuitively makes sense. But then I tested it in Snowflake and I was able to retrieve a result from the result cache using a SELECT while using a suspended warehouse (without auto resume). So A is indeed the correct answer.

upvoted 4 times

**subbie** 3 years, 9 months ago

Its true

upvoted 2 times

**SV1122** 3 years, 10 months ago

Query results are cached for 24hours and is not dependent on warehouse being up or down

upvoted 1 times

Question #134 *Topic 1*

True or False: When Snowflake is configured to use Single Sign-On (SSO), Snowflake receives the usernames and credentials from the SSO service and loads them into the customer's Snowflake account.

A. True

B. False

**Suggested Answer:** *B*

*Community vote distribution*




**Rajivnb** Highly Voted 3 years ago

No. it does not save any credentials in an Account. Purpose of Federated and SSO is to authenticate the User and pass the authorization to SSO/Snowflake. user authentication is separated from user access through the use of one or more external entities that provide independent authentication of user credentials. The authentication is then passed to one or more services, enabling users to access the services through SSO. A federated environment consists of the following components: Service provider (SP): In a Snowflake federated environment, Snowflake serves as the SP. Identity provider (IdP): The external, independent entity responsible for providing the following services to the SP: Creating and maintaining user credentials and other profile information. Authenticating users for SSO access to the SP.

upvoted 9 times

**56d43d6** Most Recent 1 year, 10 months ago

B False: https://docs.snowflake.com/en/user-guide/admin-security-fed-auth-overview#login-workflow No, the idp does NOT pass the user credentials to SF; the idp passes a SAML response to SF to authorize the initiation of a session

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B.False

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: B

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. False

upvoted 1 times

**KarBiswa** 3 years, 7 months ago

Correcting the vote to false as the credentials will be verified from IdP(Identity Provider)

upvoted 2 times

**EmiB** 3 years, 8 months ago

There's no sending credential info between SSO and SF.

upvoted 3 times

**KarBiswa** 3 years, 8 months ago

https://docs.snowflake.com/en/user-guide/admin-security-fed-auth-overview.html

upvoted 2 times

**Andywu213** 3 years, 10 months ago

False, SCIM does auto user provisioning not SSO

upvoted 1 times

**seaun** 4 years ago

Should be B

upvoted 1 times

**Himo** 4 years, 1 month ago

Yes it does. Supported SSO Workflows Federated authentication enables the following SSO workflows: Logging into Snowflake. Logging out of Snowflake. System timeout due to inactivity. The behavior for each workflow is determined by whether the action is initiated within Snowflake or your IdP. https://docs.snowflake.com/en/user-guide/admin-security-fed-auth-overview.html#supported-sso-workflowsthat means that you just logs into your organisation SSO and if you click on snowflake you dont need to login to snowflake as you are using SSO and your user data is provided - I think that is what 'loads' mean

upvoted 1 times

**Himo** 4 years, 1 month ago

**Selected Answer: A**

Yes it does. Supported SSO Workflows Federated authentication enables the following SSO workflows:

Logging into Snowflake. Logging out of Snowflake. System timeout due to inactivity.

The behavior for each workflow is determined by whether the action is initiated within Snowflake or your IdP.

https://docs.snowflake.com/en/user-guide/admin-security-fed-auth-overview.html#supported-sso-workflows

upvoted 1 times

**Himo** 4 years, 1 month ago

that means that you just logs into your organisation SSO and if you click on snowflake you dont need to login to snowflake as you are using SSO and your user data is provided - I think that is what 'loads' mean

upvoted 2 times

**Casual** 4 years, 8 months ago

B is correct

upvoted 1 times

**grabc** 4 years, 9 months ago

What do you mean by 'loads'

upvoted 2 times

Question #135 *Topic 1*

Which of the following are best practices for loading data into Snowflake? (Choose three.)

A. Aim to produce data files that are between 100 MB and 250 MB in size, compressed.

B. Load data from files in a cloud storage service in a different region or cloud platform from the service or region containing the Snowflake account, to save on cost.

C. Enclose fields that contain delimiter characters in single or double quotes.

D. Split large files into a greater number of smaller files to distribute the load among the compute resources in an active warehouse.

E. When planning which warehouse(s) to use for data loading, start with the largest warehouse possible.

F. Partition the staged data into large folders with random paths, allowing Snowflake to determine the best way to load each file.

**Suggested Answer:** *A, C, D*

*Community vote distribution*

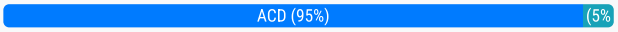


**Manocs** Highly Voted 4 years, 2 months ago

A, C, D https://docs.snowflake.com/en/user-guide/data-load-considerations.html A. To optimize the number of parallel operations for a load, we recommend aiming to produce data files roughly 100-250 MB (or larger) in size compressed. C. Fields that contain delimiter characters should be enclosed in quotes (single or double). If the data contains single or double quotes, then those quotes must be escaped. D. Split larger files into a greater number of smaller files to distribute the load among the compute resources in an active warehouse.

upvoted 10 times

**pranalig** Most Recent 2 years, 4 months ago

correct answer:ACD

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

A. Aim to produce data files that are between 100 MB and 250 MB in size, compressed. C. Enclose fields that contain delimiter characters in single or double quotes. D. Split large files into a greater number of smaller files to distribute the load among the compute resources in an active warehouse.

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: ACD

upvoted 1 times

**_yyukta** 2 years, 8 months ago

ACD are correct

upvoted 1 times

**SoftwareEngineerGoneData** 3 years, 2 months ago

Definitely ACD

upvoted 1 times

**RajuNaik** 3 years, 6 months ago

ACD right answer

upvoted 1 times

**vini15** 3 years, 6 months ago

ACD right answer

upvoted 1 times

**nishantgshah** 3 years, 7 months ago

Correct Answer ACD

upvoted 1 times

**EmiB** 3 years, 8 months ago

ACD is correct

upvoted 1 times

**Madhavan71** 3 years, 10 months ago

ACD is correct

upvoted 2 times

**baig123** 3 years, 10 months ago

correct answer ACD .

upvoted 2 times

**jveag** 4 years, 1 month ago

these tests are confusing me more than learning. the correct answers would be a, c and d.

upvoted 3 times

**William_20** 4 years, 2 months ago

b: more latency and cost, e:more cost f: less efficient

upvoted 1 times

Question #136 *Topic 1*

Which Snowflake feature is used for both querying and restoring data?

A. Cluster keys

B. Time Travel

C. Fail-safe

D. Cloning

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Time travel

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: B

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Time Travel

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Andywu213** 3 years, 10 months ago

B https://community.snowflake.com/s/article/How-to-query-time-travel-with-a-time-different-than-the-default-UTC

upvoted 1 times

**greenlever** 4 years ago

Querying, cloning, and restoring historical data in tables, schemas, and databases for up to 90 days through Snowflake Time Travel.

upvoted 4 times

**Ritab** 4 years ago

I think its an incorrect question. These two capabilities cannot be in a single feature.

upvoted 1 times

**Ritab** 4 years ago

I think its an incorrect question. These two capabilities cannot be in a single feature.Check time travel. You can query what it was like on a past date and also restore.

upvoted 1 times

**SV1122** 3 years, 10 months ago

Check time travel. You can query what it was like on a past date and also restore.

upvoted 1 times

Question #137 *Topic 1*

What do the terms scale up and scale out refer to in Snowflake? (Choose two.)

A. Scaling out adds clusters of the same size to a virtual warehouse to handle more concurrent queries.

B. Scaling out adds clusters of varying sizes to a virtual warehouse.

C. Scaling out adds additional database servers to an existing running cluster to handle more concurrent queries.

D. Snowflake recommends using both scaling up and scaling out to handle more concurrent queries.

E. Scaling up resizes a virtual warehouse so it can handle more complex workloads.

F. Scaling up adds additional database servers to an existing running cluster to handle larger workloads.

**Suggested Answer:** *A, E*

*Community vote distribution*

AE (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A and E are Correct

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: AE

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A and E are correct

upvoted 1 times

**Marge23** 3 years, 1 month ago

Answer is AE.

upvoted 2 times

**SKTumma** 3 years, 3 months ago

https://docs.snowflake.com/en/user-guide/warehouses-considerations

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**EmiB** 3 years, 7 months ago

A & E - correct

upvoted 1 times

**muks86** 3 years, 8 months ago

right ans very basic and important topic.

upvoted 1 times

**baig123** 3 years, 10 months ago

AE is correct answer.

upvoted 2 times

**Ritab** 4 years ago

The given answer sounds correct

upvoted 1 times

Question #138 *Topic 1*

What is the minimum Snowflake edition that has column-level security enabled?

A. Standard

B. Enterprise

C. Business Critical

D. Virtual Private Snowflake

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Enterprise Edition

upvoted 2 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: B

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Enterprise

upvoted 1 times

**SoftwareEngineerGoneData** 3 years, 2 months ago

https://docs.snowflake.com/en/user-guide/intro-editions

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Naveen90** 4 years, 1 month ago

https://docs.snowflake.com/en/user-guide/security-column.html#:~:text=This%20feature%20requires%20Enterprise%20Edition%20(or%20higher).

upvoted 2 times

Question #139 *Topic 1*

What parameter controls if the Virtual Warehouse starts immediately after the CREATE WAREHOUSE statement?

A. INITIALLY_SUSPENDED = TRUE/FALSE

B. START_AFTER_CREATE = TRUE/FALSE

C. START_THE = 60 // (seconds from now)

D. START_TIME = CURRENT_DATE()

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Nagesh_CL** Highly Voted 4 years ago

Answer is correct: - Syntax CREATE [ OR REPLACE ] WAREHOUSE [ IF NOT EXISTS ] [ [ WITH ] objectProperties ] [ objectParams ] Where: objectProperties ::= WAREHOUSE_SIZE = XSMALL | SMALL | MEDIUM | LARGE | XLARGE | XXLARGE | XXXLARGE | X4LARGE | X5LARGE | X6LARGE MAX_CLUSTER_COUNT = MIN_CLUSTER_COUNT = SCALING_POLICY = STANDARD | ECONOMY AUTO_SUSPEND = | NULL AUTO_RESUME = TRUE | FALSE INITIALLY_SUSPENDED = TRUE | FALSE RESOURCE_MONITOR = COMMENT = ' ' ENABLE_QUERY_ACCELERATION = TRUE | FALSE QUERY_ACCELERATION_MAX_SCALE_FACTOR =

upvoted 5 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A. INITIALLY_SUSPENDED = TRUE/FALSE

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: A

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. INITIALLY_SUSPENDED = TRUE/FALSE

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

Question #140 *Topic 1*

When cloning a database, what is cloned with the database? (Choose two.)

A. Privileges on the database

B. Existing child objects within the database

C. Future child objects within the database

D. Privileges on the schemas within the database

E. Only schemas and tables within the database

**Suggested Answer:** *B, D*

*Community vote distribution*

BD (100%)

**SandipMore** Highly Voted 3 years, 2 months ago

A. Privileges on the database -- This is not correct as source for cloning do not inherit privileges B. Existing child objects within the database --- This is correct C. Future child objects within the database --- This is not correct D. Privileges on the schemas within the database Most Voted -- As Schema is child, it inherits all grants E. Only schemas and tables within the database --- This is not correct

upvoted 5 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A and B are Correct

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: BD

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B and D are correct

upvoted 1 times

**Roohee** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/object-clone

upvoted 1 times

**pri060** 3 years, 2 months ago

only B is correct. Grants are only inherited if Copy Grants keyword is used while creating a clone.

upvoted 4 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 2 times

**RajuNaik** 3 years, 6 months ago

BD is correct answer

upvoted 1 times

**SV1122** 3 years, 10 months ago

BD. If the source object is a database or schema, the clone inherits all granted privileges on the clones of all child objects contained in the source object. https://docs.snowflake.com/en/user-guide/object-clone.html#cloning-and-governance-objects

upvoted 2 times

**aalonso** 3 years, 10 months ago

child objects grants are cloned

upvoted 2 times

**sdnv89** 3 years, 11 months ago

Database itself plus all child objects. Not the grants on the database itself, but the grants on the child objects are preserved.
https://docs.snowflake.com/en/user-guide/object-clone.html#access-control-privileges-for-cloned-objects

upvoted 3 times

**William_20** 4 years, 2 months ago

All objects and grants are cloned except on the clone level (in this case database)

upvoted 1 times

**Manocs** 4 years, 2 months ago

B, D For D: If the source object is a database or schema, for child objects contained in the source, the clone replicates all granted privileges on the corresponding child objects https://docs.snowflake.com/en/user-guide/object-clone.html#access-control-privileges-for-cloned-objects

upvoted 3 times

Question #141 *Topic 1*

Which of the following describes the Snowflake Cloud Services layer?

A. Coordinates activities in the Snowflake account

B. Executes queries submitted by the Snowflake account users

C. Manages quotas on the Snowflake account storage

D. Manages the virtual warehouse cache to speed up queries

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**EmiB** Highly Voted 3 years, 7 months ago

A. Coordinates activities in the Snowflake account - CLOUD SERVICE LAYER B. Executes queries submitted by the Snowflake account users - COMPUTE LAYER C. Manages quotas on the Snowflake account storage - STORAGE LAYER D. Manages the virtual warehouse cache to speed up queries - COMPUTE LAYERad C. - there is no STORAGE QUATA, only CREDIT QUATA

upvoted 17 times

**EmiB** Highly Voted 3 years, 7 months ago

Selected Answer: A

A. Coordinates activities in the Snowflake account - CLOUD SERVICE LAYER B. Executes queries submitted by the Snowflake account users - COMPUTE LAYER C. Manages quotas on the Snowflake account storage - STORAGE LAYER D. Manages the virtual warehouse cache to speed up queries - COMPUTE LAYER

upvoted 17 times

**EmiB** 3 years, 7 months ago

ad C. - there is no STORAGE QUATA, only CREDIT QUATA

upvoted 4 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

It coordinates activities in the snowflake,A is correct Answer

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: A

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. Coordinates activities in the Snowflake account

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**JohnPi** 3 years, 12 months ago

The cloud services layer is a collection of services that coordinate activities across Snowflake

upvoted 4 times

**Penn_W** 4 years, 1 month ago

A is correct. The cloud services layer is a set of services that help Snowflake coordinate activities.

upvoted 3 times

**stamford** 4 years, 1 month ago

Answer is D

upvoted 1 times

Question #142 *Topic 1*

What is the maximum total Continuous Data Protection (CDP) charges incurred for a temporary table?

A. 30 days

B. 7 days

C. 48 hours

D. 24 hours

**Suggested Answer:** *D*

*Community vote distribution*




**visionary4ever** Most Recent 1 year, 2 months ago

Default = 1 day (24 hrs). Maximum = 48 hrs.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

D.24 hrs

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: D

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. 24 hours

upvoted 1 times

**SoftwareEngineerGoneData** 3 years, 2 months ago

https://docs.snowflake.com/en/user-guide/data-cdp-storage-costs

upvoted 3 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**kjdfsafc87456** 4 years, 1 month ago

Thus, the maximum total CDP charges incurred for a temporary table are 1 day (or less if the table is explicitly dropped or dropped as a result of terminating the session). During this period, Time Travel can be performed on the table. https://docs.snowflake.com/en/user-guide/tables-storage-considerations.html

upvoted 2 times

Question #143 *Topic 1*

When reviewing a query profile, what is a symptom that a query is too large to fit into the memory?

A. A single join node uses more than 50% of the query time

B. Partitions scanned is equal to partitions total

C. An AggregateOperator node is present

D. The query is spilling to remote storage

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**omiiUST** Most Recent 2 years, 3 months ago

"Bytes spilled to remote storage" indicates the volume of data that couldn't fit in memory and had to be spilled to a virtual warehouse temporary storage.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

D.is correct

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct Answer: D

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: D

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D is correct

upvoted 1 times

**RajuNaik** 3 years, 6 months ago

Answer should be D

upvoted 2 times

**c0d3g** 3 years, 7 months ago

Selected Answer: D

For some operations (e.g. duplicate elimination for a huge data set), the amount of memory available for the compute resources used to execute the operation might not be sufficient to hold intermediate results. As a result, the query processing engine will start *spilling* the data to local disk. If the local disk space is not sufficient, the spilled data is then saved to remote disks. https://docs.snowflake.com/en/user-guide/ui-query-profile.html#queries-too-large-to-fit-in-memory

upvoted 4 times

**c0d3g** 3 years, 7 months ago

For some operations (e.g. duplicate elimination for a huge data set), the amount of memory available for the compute resources used to execute the operation might not be sufficient to hold intermediate results. As a result, the query processing engine will start *spilling* the data to local disk. If the local disk space is not sufficient, the spilled data is then saved to remote disks. https://docs.snowflake.com/en/user-guide/ui-query-profile.html#queries-too-large-to-fit-in-memoryI agree, I think having all partitions scanned just means you might want to set up a cluster key.

upvoted 4 times

**whiteomax** 2 years, 9 months ago

I agree, I think having all partitions scanned just means you might want to set up a cluster key.

upvoted 2 times

**EmiB** 3 years, 8 months ago

As a result, the query processing engine will start spilling the data to local disk. If the local disk space is not sufficient, the spilled data is then saved to remote disks.

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

**Selected Answer: D**

as per documentaion

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

as per documentaionas per the document it will spill to local storage and if it fills then dumped into Remote storage

upvoted 1 times

**JRTanubuddi** 3 years, 8 months ago

as per the document it will spill to local storage and if it fills then dumped into Remote storage

upvoted 1 times

**iamPhD** 3 years, 10 months ago

Answer should be D

upvoted 1 times

**aalonso** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/ui-query-profile.html#queries-too-large-to-fit-in-memory

upvoted 2 times

**baig123** 3 years, 10 months ago

I agree with Answer D, if query spilling happens , query processing engine start saving the data to local disk if this is not sufficient then it start saving the data to remote disk, this process significantly impact query performance, the solution is use larger size warehouse.

upvoted 1 times

**kjdfsafc87456** 4 years, 2 months ago

See docs referred to by Manocs.

upvoted 1 times

**William_20** 4 years, 2 months ago

See docs referred to by Manocs.

upvoted 1 times

**Manocs** 4 years, 2 months ago

https://docs.snowflake.com/en/user-guide/ui-query-profile.html#queries-too-large-to-fit-in-memory

upvoted 1 times

Question #144 *Topic 1*

What type of query benefits the MOST from search optimization?

A. A query that uses only disjunction (i.e., OR) predicates

B. A query that includes analytical expressions

C. A query that uses equality predicates or predicates that use IN

D. A query that filters on semi-structured data types

**Suggested Answer:** *C*

*Community vote distribution*

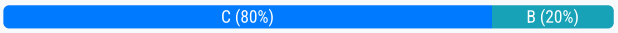


**SV1122** Highly Voted 3 years, 10 months ago

The correct answer is C. 2: Below are the recommended checks to consider for the query: At least one of the columns accessed through the query filter operation has at least 100k-200k distinct values. The query uses equality predicate or predicates that use IN. The query returns a few rows with highly selective filters. The query typically runs for at least tens of seconds. https://community.snowflake.com/s/article/Search-Optimization-When-How-To-Use

upvoted 8 times

**SV1122** 4 months ago

Was on exam Dec 31st 2022

upvoted 3 times

**Mallikharjuna452** 2 years, 3 months ago

C.A query that uses equality predicates or predicates that use IN

upvoted 1 times

**EmiB** 4 months ago

Selected Answer: C, D

Using a query as the source for the COPY statement for column reordering [C], column omission [D], and casts (i.e. transforming data during a load) is supported. For usage examples, see Transforming Data During a Load. Note that only simple SELECT statements are supported. Filtering using a WHERE clause is not supported [E].

upvoted 2 times

**pranalig** 2 years, 4 months ago

Correct Answer: C

upvoted 1 times

**SV1122** 4 months ago

CD is correct https://docs.snowflake.com/en/sql-reference/sql/create-pipe.html

upvoted 2 times

**_yyukta** 2 years, 8 months ago

C. A query that uses equality predicates or predicates that use IN

upvoted 1 times

**Mallikharjuna452** 4 months ago

C and D are correct

upvoted 1 times

**BobFar** 3 years ago

Correct Answer is B. Analytical expression , the answer C is only for Variant, Objects and Array data type. https://docs.snowflake.com/en/user-guide/search-optimization-serviceNot correct. Correct answer is CC is correct. The service works on the predicate.

upvoted 1 times

**pranalig** Most Recent 4 months ago

Correct Answer: CD

upvoted 1 times

**Learner21525** 3 years ago

Not correct. Correct answer is C

upvoted 2 times

**Mike722** 2 years, 8 months ago

C is correct. The service works on the predicate.

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**PoornaChandraIn** 3 years, 3 months ago

The answer is C.

upvoted 1 times

**junpal** 3 years, 6 months ago

C. A query that uses equality predicates or predicates that use IN

upvoted 2 times

**VisionBI** 3 years, 7 months ago

Should be B. The search optimization service can significantly improve the performance of certain types of lookup and analytical queries that use an extensive set of predicates for filtering. Query uses equality predicate or predicates that use IN only applies to VARIANT, OBJECT, and ARRAY columns . https://docs.snowflake.com/en/user-guide/search-optimization-serviceB is false: "certain types of ... analytical queries"

upvoted 1 times

**Def21** 3 years, 4 months ago

B is false: "certain types of ... analytical queries"

upvoted 1 times

Question #145 *Topic 1*

What transformations are supported in a CREATE PIPE ... AS COPY `¦ FROM (`¦) statement? (Choose two.)

A. Data can be filtered by an optional WHERE clause.

B. Incoming data can be joined with other tables.

C. Columns can be reordered.

D. Columns can be omitted.

E. Row level access can be defined.

**Suggested Answer:** *C, D*

*Community vote distribution*

CD (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

C and D are correct

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct Answer: CD

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: CD

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C and D are correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**EmiB** 3 years, 8 months ago

Using a query as the source for the COPY statement for column reordering [C], column omission [D], and casts (i.e. transforming data during a load) is supported. For usage examples, see Transforming Data During a Load. Note that only simple SELECT statements are supported. Filtering using a WHERE clause is not supported [E].

upvoted 3 times

**KarBiswa** 3 years, 8 months ago

as per documentation

upvoted 1 times

**SV1122** 3 years, 10 months ago

Was on exam Dec 31st 2022

upvoted 3 times

**SV1122** 3 years, 10 months ago

CD is correct https://docs.snowflake.com/en/sql-reference/sql/create-pipe.html

upvoted 2 times

**SV1122** 3 years, 10 months ago

CD is correct https://docs.snowflake.com/en/sql-reference/sql/create-pipe.htmlhttps://docs.snowflake.com/en/sql-reference/sql/create-pipe.html#usage-noteshttps://docs.snowflake.com/en/user-guide/data-load-transform.html

upvoted 2 times

**BigDataBB** 3 years, 10 months ago

https://docs.snowflake.com/en/sql-reference/sql/create-pipe.html#usage-notes

upvoted 1 times

**BigDataBB** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/data-load-transform.html

upvoted 1 times

Question #146 *Topic 1*

Which of the following are characteristics of Snowflake virtual warehouses? (Choose two.)

A. Auto-resume applies only to the last warehouse that was started in a multi-cluster warehouse.

B. The ability to auto-suspend a warehouse is only available in the Enterprise edition or above.

C. SnowSQL supports both a configuration file and a command line option for specifying a default warehouse.

D. A user cannot specify a default warehouse when using the ODBC driver.

E. The default virtual warehouse size can be changed at any time.

**Suggested Answer:** *C*

*Community vote distribution*




**nintendogamer64** Highly Voted 3 years, 10 months ago

CE are correct. I cant believe exam topics wants to charge for this test. We cant even vote for the correct answer in this case. What a joke.

upvoted 32 times

**alexhparreira** Most Recent 5 months ago

configuration file and a command line option for specifying a default warehouse.: In SnowSQL (Snowflake's command-line client), you can specify your target virtual warehouse dynamically using the -w or --warehouse flag during connection, or save it permanently as a parameter within the [options] section of the standard config file. E. The default virtual warehouse size can be changed at any time.: Administrators or authorized roles can scale a virtual warehouse size up or down instantly (e.g., from X-Small to Large) at any point using an ALTER WAREHOUSE DDL command. This can even be done dynamically while the warehouse is actively running and executing queries without needing a restart.

upvoted 1 times

**Stockbug** 11 months ago

CE but somehow examtopics adjusted only one selection answer.

upvoted 2 times

**MultiCloudIronMan** 1 year, 10 months ago

CE is correct

upvoted 2 times

**Kuldeepkd** 1 year, 11 months ago

C and E both

upvoted 2 times

**Shipra123** 1 year, 11 months ago

Correct answer C and E.

upvoted 1 times

**Daniel1412** 2 years ago

C and E

upvoted 2 times

**Alongi2** 2 years ago

C and E

upvoted 1 times

**ABRocks** 2 years, 2 months ago

C and E are correct

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

C and E are correct

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct Answer: CE

upvoted 2 times

**gunasekarantjk1996** 2 years, 5 months ago

CE. Still exam topics we cant blindly believe.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

CE correct

upvoted 1 times

**sandy91** 2 years, 9 months ago

C and E

upvoted 1 times

**Colourseun** 2 years, 12 months ago

C and E Are correct

upvoted 1 times

**Heetec** 3 years ago

**Selected Answer: E**

CE correct

upvoted 2 times

**Afzy** 3 years ago

C AND E

upvoted 2 times

Question #147 *Topic 1*

Which command should be used to load data from a file, located in an external stage, into a table in Snowflake?

A. INSERT

B. PUT

C. GET

D. COPY

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**56d43d6** Most Recent 1 year, 10 months ago

https://docs.snowflake.com/en/sql-reference/sql/copy-into-table

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

COPY COMMAND

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct Answer: D

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. COPY

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**EmiB** 3 years, 7 months ago

D - COPY INTO exactly

upvoted 1 times

**NabilR** 3 years, 9 months ago

Why not INSERT ?

upvoted 1 times

**SV1122** 3 years, 10 months ago

GET does not support downloading files from external stages https://docs.snowflake.com/en/sql-reference/sql/get.html So the answer is COPY

upvoted 1 times

**sdnv89** 3 years, 11 months ago

copy into table

upvoted 1 times

**MMT123764k** 3 years, 12 months ago

Correct answer Copy

upvoted 1 times

Question #148 *Topic 1*

The Snowflake Cloud Data Platform is described as having which of the following architectures?

A. Shared-disk

B. Shared-nothing

C. Multi-cluster shared data

D. Serverless query engine

**Suggested Answer:** *C*

*Community vote distribution*

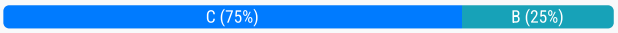


**Mallikharjuna452** Most Recent 2 years, 3 months ago

C.Multi clustered shared data

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct Answer: C

upvoted 1 times

**MBrat** 2 years, 5 months ago

https://docs.snowflake.com/en/user-guide/intro-key-concepts#snowflake-architecture "Snowflake uses a central data repository for persisted data that is accessible from all compute nodes in the platform" So it is definitely shared data "Snowflake processes queries using MPP (massively parallel processing) compute clusters" And it uses multiple clusters for processing. So the answer is C, Multi-cliuster shared data. Sometimes the question answers aren't verbatim from the docs, you have to understand the meaning of words.

upvoted 2 times

**Alex_Ovidiu** 2 years, 7 months ago

The Snowflake Cloud Data Platform is described as having a D. Serverless query engine architecture. While Snowflake leverages elements of both shared-disk and shared-nothing architectures for data storage and processing, the core aspect that defines its architecture is the serverless nature of the query engine.

upvoted 1 times

**rogerek077** 2 years, 7 months ago

https://docs.snowflake.com/en/user-guide/intro-key-concepts#snowflake-architecture "Snowflake's architecture is a hybrid of traditional shared-disk and shared-nothing database architectures." On the day of March 28th 2024 I could not find any explicit mention of Multi-cluster shared data architecture in the mentioned article. I guess this question is quite obsolete or maybe the most commonly voted answer is a kind of a synonym for "a hybrid of traditional shared-disk and shared-nothing". Not fully sure

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C. Multi-cluster shared data

upvoted 1 times

**VenkataPhaniPavanKumar** 2 years, 9 months ago

Correct Answer : C The Snowflake data platform has multi-cluster shared data architecture is designed to handle large data volumes with speed and efficiency. All data processing horsepower within Snowflake is managed by one or more clusters of computing resources.

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 2 times

**learner2023** 3 years, 6 months ago

Snowflake's architecture is a hybrid of traditional shared-disk and shared-nothing database architectures. Similar to shared-disk architectures, Snowflake uses a central data repository for persisted data that is accessible from all compute nodes in the platform.

upvoted 3 times

**EmiB** 3 years, 7 months ago

**Selected Answer: C**

C. Multi-cluster shared data

upvoted 3 times

**EmiB** 3 years, 7 months ago

C. Multi-cluster shared dataThis might be true but this is not a typical description: "is described as"

upvoted 3 times

**Def21** 3 years, 4 months ago

This might be true but this is not a typical description: "is described as"

upvoted 1 times

**AlexbDku** 3 years, 10 months ago

C - multiple clusters access the same data

upvoted 1 times

**SV1122** 3 years, 10 months ago

Correct answer is C - The storage is shared however the compute is multi clustered https://www.snowflake.com/product/architecture/

upvoted 3 times

**baig123** 3 years, 10 months ago

correct answer is A and B , Snowflake's architecture is a hybrid of traditional shared-disk and shared-nothing database architectures

upvoted 3 times

Question #149 *Topic 1*

Which of the following is a data tokenization integration partner?

A. Protegrity

B. Tableau

C. DBeaver

D. SAP

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**kjdfsafc87456** Highly Voted 4 years, 2 months ago

The following partners facilitate external tokenization in Snowflake. To use these partner integrations, follow the instructions in the partner documentation or contact the partner to begin the configuration process: ALTR Baffle Fortanix MicroFocus CyberRes Voltage Protegrity Privacera SecuPI Skyflow

upvoted 7 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A.Protegrity

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. Protegrity

upvoted 1 times

**Arnab14** 3 years, 5 months ago

ALTR Baffle Comforte Fortanix MicroFocus CyberRes Voltage Protegrity Privacera SecuPI Skyflow Spring Labs https://docs.snowflake.com/en/user-guide/security-column-ext-token-use

upvoted 4 times

**dmitriypo** 3 years, 6 months ago

Protegrity: https://docs.snowflake.com/en/user-guide/security-column-ext-token-use

upvoted 1 times

Question #150 *Topic 1*

What versions of Snowflake should be used to manage compliance with Personal Identifiable Information (PII) requirements? (Choose two.)

A. Custom Edition

B. Virtual Private Snowflake

C. Business Critical Edition

D. Standard Edition

E. Enterprise Edition

**Suggested Answer:** *B, C*

*Community vote distribution*




**seaun** Highly Voted 4 years ago

Answer should be BC https://docs.snowflake.com/en/user-guide/intro-editions.html

upvoted 7 times

**PrasantSadangi** 1 year, 7 months ago

C & E https://docs.snowflake.com/en/user-guide/classify-intro https://docs.snowflake.com/en/user-guide/classify-customNo, According to the Snowflake documentation: Business Critical Edition (formerly Enterprise for Sensitive Data) is specifically designed for organizations that handle highly sensitive data, including PHI/PII and workloads that must comply with regulations such as HIPAA and HITRUST. It adds enhanced security and data protection capabilities beyond Enterprise Edition. (docs.snowflake.com) Virtual Private Snowflake (VPS) provides Snowflake's highest level of security. It includes all Business Critical features and runs in a completely isolated Snowflake environment, making it suitable for organizations with the strictest data protection and compliance requirements. (docs.snowflake.com) Enterprise Edition  Adds enterprise features, but Snowflake identifies Business Critical as the edition intended for sensitive data compliance requirements such as HIPAA

upvoted 1 times

**PrasantSadangi** 1 year, 7 months ago

Selected Answer: CE

C & E

https://docs.snowflake.com/en/user-guide/classify-intro https://docs.snowflake.com/en/user-guide/classify-custom

upvoted 1 times

**BigDataBB** Most Recent 4 months, 1 week ago

No, According to the Snowflake documentation: Business Critical Edition (formerly Enterprise for Sensitive Data) is specifically designed for organizations that handle highly sensitive data, including PHI/PII and workloads that must comply with regulations such as HIPAA and HITRUST. It adds enhanced security and data protection capabilities beyond Enterprise Edition. (docs.snowflake.com) Virtual Private Snowflake (VPS) provides Snowflake's highest level of security. It includes all Business Critical features and runs in a completely isolated Snowflake environment, making it suitable for organizations with the strictest data protection and compliance requirements. (docs.snowflake.com) Enterprise Edition  Adds enterprise features, but Snowflake identifies Business Critical as the edition intended for sensitive data compliance requirements such as HIPAA

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct Answer: BC

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: BC

upvoted 1 times

**matias_salina** 3 years ago

**Selected Answer: BC**

"Business Critical Edition, formerly known as Enterprise for Sensitive Data (ESD)'"

https://docs.snowflake.com/en/user-guide/intro-editions#business-critical-edition https://community.snowflake.com/s/article/6-Steps-to-Secure-PII-in-Snowflake-s-Cloud-Data-Platform

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B and C

upvoted 1 times

**matias_salina** 3 years ago

"Business Critical Edition, formerly known as Enterprise for Sensitive Data (ESD)'" https://docs.snowflake.com/en/user-guide/intro-editions#business-critical-edition https://community.snowflake.com/s/article/6-Steps-to-Secure-PII-in-Snowflake-s-Cloud-Data-PlatformIf E is correct, means that B and C are correct too, so B and C are for sure correct. The approach mentioned in one of the links above seems to be achievable with the standard edition as well, as it uses secure views and data masking (both features are available in the standard edition)

upvoted 1 times

**matias_salina** 3 years ago

If E is correct, means that B and C are correct too, so B and C are for sure correct. The approach mentioned in one of the links above seems to be achievable with the standard edition as well, as it uses secure views and data masking (both features are available in the standard edition)

upvoted 1 times

**otsumy** 3 years, 2 months ago

C and E. Snowflake Enterprise for Sensitive Data (ESD) provides a feature set which allows its users to easily protect sensitive PII, even allowing the possibility of meeting HIPAA regulations where required. Less stringent protections may be met using only the standard Snowflake DB offering. https://community.snowflake.com/s/article/6-Steps-to-Secure-PII-in-Snowflake-s-Cloud-Data-Platform

upvoted 1 times

**DataEngineer7331** 3 years, 6 months ago

This might be a tricky one. In my opinion, Data Masking is in compliance with PII. So the Enterprise Version would also be in compliance with PII. There is also Column Level Security. BUT: VPS also features the same... https://docs.snowflake.com/en/user-guide/governance-classify-concepts.html and https://docs.snowflake.com/en/user-guide/intro-editions.html

upvoted 2 times

**KarBiswa** 3 years, 8 months ago

BC correct

upvoted 1 times

**iamPhD** 3 years, 10 months ago

BC are correct

upvoted 1 times

**mbacelar** 4 years, 1 month ago

**Selected Answer: CE**

https://docs.snowflake.com/en/user-guide/governance-classify-concepts.html This feature requires Enterprise Edition or higher. There is no Virtual Private Snowflake

upvoted 1 times

**seaun** 4 years ago

Answer should be BC

upvoted 3 times

**mbacelar** 4 years, 1 month ago

https://docs.snowflake.com/en/user-guide/governance-classify-concepts.html This feature requires Enterprise Edition or higher. There is no Virtual Private SnowflakeThere is Virtual Private Snowflake. Check here https://docs.snowflake.com/en/user-guide/intro-editions.html

upvoted 1 times

**SV1122** 3 years, 10 months ago

There is Virtual Private Snowflake. Check here https://docs.snowflake.com/en/user-guide/intro-editions.html

upvoted 1 times

Question #151 *Topic 1*

What are supported file formats for unloading data from Snowflake? (Choose three.)

A. XML

B. JSON

C. Parquet

D. ORC

E. AVRO

F. CSV

**Suggested Answer:** *B, C, F*

*Community vote distribution*

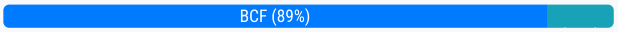


**IlaS** Most Recent 10 months, 3 weeks ago

All of them are now supported as per the snowflake documentation.

upvoted 1 times

**gunasekarantjk1996** 1 year, 11 months ago

File formats Delimited files (CSV, TSV, etc.) Any valid delimiter is supported; default is comma (i.e. CSV). JSON Parquet

upvoted 2 times

**AnireddySaikiranReddy** 2 years, 2 months ago

Correct Answer: BCF

upvoted 1 times

**_yyukta** 2 years, 2 months ago

BCF are correct

upvoted 2 times

**PoornaChandraIn** 2 years, 9 months ago

XML for loading data, ORC, AVRO for un loading semi structed data. CSV, JSON, Parquet for unloading structed data . So BCD are the correct answer.

upvoted 1 times

**c0d3g** 3 years, 1 month ago

JSON Parquet CSV https://docs.snowflake.com/en/user-guide/data-unload-prepare

upvoted 2 times

**mayank** 3 years, 3 months ago

even snowflake supports XML. Why this is not the correct answer.

upvoted 2 times

**mayank** 3 years, 3 months ago

I think JSON,CSV and XML should be the right answer

upvoted 1 times

**mayank** 3 years, 3 months ago

even snowflake supports XML. Why this is not the correct answer.XML is only supported for loading, not unloading

upvoted 2 times

**VisionBI** 3 years ago

XML is only supported for loading, not unloading

upvoted 3 times

**iamPhD** 3 years, 4 months ago

B,C,F are correct

upvoted 2 times

**baig123** 3 years, 4 months ago

BC and F is correct answer.

upvoted 1 times

**greenlever** 3 years, 6 months ago

correct

upvoted 2 times

Question #152 *Topic 1*

The Snowflake cloud services layer is responsible for which tasks? (Choose two.)

A. Local disk caching

B. Authentication and access control

C. Metadata management

D. Query processing

E. Database storage

**Suggested Answer:** *B, C*

*Community vote distribution*

BC (100%)

**a8c5bbe** Most Recent 1 year, 11 months ago

B and C correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B and C are Correct

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct Answer: BC

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: CD

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B and C are correct

upvoted 1 times

**singhks** 3 years, 2 months ago

B and C

upvoted 1 times

**isolated_sniper** 3 years, 6 months ago

Correct answer should be BC Cloud Services: The cloud services layer is a collection of services that coordinate activities across Snowflake. These services tie together all of the different components of Snowflake in order to process user requests, from login to query dispatch. The cloud services layer also runs on compute instances provisioned by Snowflake from the cloud provider. Services managed in this layer include: Authentication Infrastructure management Metadata management Query parsing and optimization Access control

upvoted 2 times

**RajuNaik** 3 years, 6 months ago

B and C are correct answer.

upvoted 1 times

**c0d3g** 3 years, 7 months ago

Authentication and access control Metadata management

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

https://docs.snowflake.com/en/user-guide/intro-key-concepts.html

upvoted 1 times

**baig123** 3 years, 10 months ago

B and C are correct answer.

upvoted 2 times

**sdnv89** 3 years, 11 months ago

B and C are correct

upvoted 2 times

**stornati** 3 years, 11 months ago

Clearly explained in the figure in the arch diagram. https://docs.snowflake.com/en/user-guide/intro-key-concepts.html

upvoted 2 times

**seaun** 4 years ago

Answer should be BC

upvoted 1 times

**greenlever** 4 years ago

correct

upvoted 1 times

**Naveen90** 4 years, 1 month ago

B and C are the correct answers

upvoted 1 times

**jveag** 4 years, 1 month ago

I think B and C.

upvoted 1 times

Question #153 *Topic 1*

What is a key feature of Snowflake architecture?

A. Zero-copy cloning creates a mirror copy of a database that updates with the original.

B. Software updates are automatically applied on a quarterly basis.

C. Snowflake eliminates resource contention with its virtual warehouse implementation.

D. Multi-cluster warehouses allow users to run a query that spans across multiple clusters.

E. Snowflake automatically sorts DATE columns during ingest, for fast retrieval by date.

**Suggested Answer:** *C*

*Community vote distribution*




**56d43d6** Highly Voted 1 year, 10 months ago

D is wrong. A single query does not span multiple clusters. It spans multiple nodes in a single cluster. This is why you expand the size of the WH to accommodate larger queries. C is correct.

upvoted 5 times

**Altria** 1 year, 2 months ago

C. Snowflake eliminates resource contention with its virtual warehouse implementation. C is correct

upvoted 2 times

**Alongi2** 2 years ago

The main feature

upvoted 1 times

**VLMBL_94** 2 years, 2 months ago

The answer is C. Multi cluster VWH supports concurrency meaning multiple queries can run in multiple clusters not one query spanning across multiple clusters

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

D Is correct,Multi Culster Warehouses contention with its virtual warehouse implementation.

upvoted 1 times

**pranalig** 2 years, 4 months ago

D is correct. Multi cluster warehouses is the unique feature of SF.

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

C. Snowflake eliminates resource contention with its virtual warehouse implementation. Snowflake's architecture is built around the concept of separating compute resources from storage, allowing for scalability and eliminating resource contention through its virtual warehouse implementation. This architecture enables users to scale their compute resources up or down as needed without impacting the underlying data storage.

upvoted 1 times

**junaid2107** 2 years, 11 months ago

D based on docs

upvoted 4 times

**Colourseun** 2 years, 11 months ago

A. Zero-copy cloning creates a mirror copy of a database that updates with the original. is a key feature of Snowflake architecture. Snowflake's zero-copy cloning feature creates a snapshot of a database that is always up-to-date with the original database. This allows users to create clones of their databases for development, testing, and disaster recovery purposes without having to copy any data. Zero-copy cloning is also used to create time travel snapshots of databases.No it's not updted , any DL operation perfomraed on Original table hasve no effetc on CLined table ND VICE VERSA

upvoted 1 times

**Colourseun** 2 years, 11 months ago

A. Zero-copy cloning creates a mirror copy of a database that updates with the original. is a key feature of Snowflake architecture.

Snowflake's zero-copy cloning feature creates a snapshot of a database that is always up-to-date with the original database. This allows users to create clones of their databases for development, testing, and disaster recovery purposes without having to copy any data. Zero-copy cloning is also used to create time travel snapshots of databases.

upvoted 1 times

**afeafa** 1 year, 10 months ago

No it's not updted , any DL operation perfomraed on Original table hasve no effetc on CLined table ND VICE VERSA

upvoted 1 times

**matias_salina** 3 years, 1 month ago

Snowflake eliminates resource contention with multi-cluster " The overarching performance challenge with traditional data warehouses comes down to one thing: resource contention. Single-cluster systems—both in shared-disk and shared-nothing architectures—require all users in the data warehouse to share a limited set of resources. Too many concurrent users or workloads can lead to significant performance degradation and queueing. In contrast, Snowflake provides a patented multicluster, shared data architecture and intelligent use of readily available cloud resources to allow essentially limitless concurrency..." https://www.snowflake.com/wp-content/uploads/2020/01/how-snowflake-automates-performance-in-a-modern-cloud-data-warehouse.pdf#:~:text=Dedicated%20and%20separate%20virtual%20warehouses,It%20also%20eliminates%20contention.No, multicluster is to manage more user, a SINGLE query is NOT distributed in more cluster of the WH

upvoted 3 times

**matias_salina** 3 years, 1 month ago

**Selected Answer: D**

Snowflake eliminates resource contention with multi-cluster

" The overarching performance challenge with traditional data warehouses comes down to one thing: resource contention. Single-cluster systems—both in shared-disk and shared-nothing architectures—require all users in the data warehouse to share a limited set of resources. Too many concurrent users or workloads can lead to significant performance degradation and queueing. In contrast, Snowflake provides a patented multicluster, shared data architecture and intelligent use of readily available cloud resources to allow essentially limitless concurrency..."

https://www.snowflake.com/wp-content/uploads/2020/01/how-snowflake-automates-performance-in-a-modern-cloud-data-warehouse.pdf#:~:text=Dedicated%20and%20separate%20virtual%20warehouses,It%20also%20eliminates%20contention.

upvoted 3 times

**BigDataBB** Most Recent 4 months, 1 week ago

No, multicluster is to manage more user, a SINGLE query is NOT distributed in more cluster of the WH

upvoted 1 times

**Marge23** 3 years, 1 month ago

D is correct. Multi cluster warehouses is the unique feature of SF.

upvoted 1 times

**SkipD** 3 years, 6 months ago

"eliminates" is a strong word though, no? That skews me toward D.....

upvoted 1 times

**c0d3g** 3 years, 7 months ago

Snowflake eliminates resource contention with its virtual warehouse implementation.

upvoted 1 times

**Raju039** 3 years, 7 months ago

C is correct

upvoted 1 times

**fahfouhi94** 3 years, 8 months ago

D is also correct but qst is about the key feature of snowflake

upvoted 1 times

**baig123** 3 years, 10 months ago

C and D both correct answers

upvoted 2 times

**KarBiswa** 3 years, 8 months ago

if its a single selection

upvoted 1 times

**baig123** 3 years, 10 months ago

C and D both correct answersDifferent queries goes in different cluster a single query is processed in a single cluster. Multi cluster is for "Concurrency"

upvoted 2 times

**BigDataBB** Highly Voted 3 years, 10 months ago

Different queries goes in different cluster a single query is processed in a single cluster. Multi cluster is for "Concurrency"

upvoted 8 times

Question #154 *Topic 1*

When publishing a Snowflake Data Marketplace listing into a remote region what should be taken into consideration? (Choose two.)

A. There is no need to have a Snowflake account in the target region, a share will be created for each user.

B. The listing is replicated into all selected regions automatically, the data is not.

C. The user must have the ORGADMIN role available in at least one account to link accounts for replication.

D. Shares attached to listings in remote regions can be viewed from any account in an organization.

E. For a standard listing the user can wait until the first customer requests the data before replicating it to the target region.

**Suggested Answer:** *B, E*

*Community vote distribution*




**scoob** Highly Voted 3 years, 4 months ago

BE A - False - To share data in a region, you must have an account in that region in order to replicate data. If you have more than one account, all accounts must belong to the same organization B - True - To share data in a region, you must have an account in that region in order to replicate data. If you have more than one account, all accounts must belong to the same organization C - True, but not relevant to Data Marketplace D - false - When you publish a listing in a remote region, you can either allow all accounts in your organization to fulfill listing requests or explicitly add individual accounts as providers. Only the listing owner can specify who can fulfill listing requests. E - True - For free listings, you must replicate data to each of the selected regions before publishing the listing. For standard listings, you do not need to replicate the data to each region until a consumer requests it. https://docs.snowflake.com/en/user-guide/data-exchange-managing-data-listings.html#considerations-for-creating-a-listing-in-a-remote-region-and-replicating-data

upvoted 22 times

**KarBiswa** Highly Voted 3 years, 2 months ago

Agreed from https://docs.snowflake.com/en/user-guide/data-exchange-managing-data-listings.html#considerations-for-creating-a-listing-in-a-remote-region-and-replicating-data

upvoted 7 times

**_yyukta** Most Recent 2 years, 2 months ago

B and e are correct

upvoted 1 times

**0e504b5** 2 years, 3 months ago

https://docs.snowflake.com/en/user-guide/data-exchange-managing-data-listings While listings are automatically replicated, the data is not. You do not need to replicate the data to each region until a consumer requests it.

upvoted 1 times

**singhks** 2 years, 8 months ago

https://other-docs.snowflake.com/en/collaboration/provider-listings-auto-fulfillment Configuring Cross-Cloud Auto-fulfillment If you want to automatically replicate the data product associated with your listing to other Snowflake regions, configure Cross-Cloud Auto-fulfillment. When auto-fulfillment is enabled for a listing, Snowflake automatically replicates your data product to consumer regions as needed. The data product is the tables, schemas, UDFs, UDTFs, views, etc. that are part of your listing. By using auto-fulfillment, you can avoid manually replicating your data products and approving requests for your listings, helping consumers access your listings faster. When you publish a listing, it is visible to consumers in all current and future Marketplace regions, but consumers can only get, purchase, or request your product in regions you select. See Configuring Cross-Cloud Auto-fulfillment for more about region availability.

upvoted 1 times

**mlaus** 3 years, 3 months ago

B, E ist right

upvoted 3 times

**db1290** 3 years, 5 months ago

B is certain, C is the next most probable.

upvoted 1 times

**William_20** 3 years, 8 months ago

B, C. B is right (https://docs.snowflake.com/en/user-guide/data-exchange-managing-data-listings.html#considerations-for-creating-a-listing-in-a-remote-region-and-replicating-data). C because of deduction. But makes sense. You need ORGADMIN to manage multiple accounts.I think its BE: No, you need ACCOUNTADMIN role, or have the OWNERSHIP privilege. https://other-docs.snowflake.com/en/collaboration/provider-listings-creating-publishing.html#create-a-listing-for-the-sf-marketplace

upvoted 2 times

**William_20** 3 years, 8 months ago

**Selected Answer: BC**

B, C. B is right (https://docs.snowflake.com/en/user-guide/data-exchange-managing-data-listings.html#considerations-for-creating-a-listing-in-a-remote-region-and-replicating-data). C because of deduction. But makes sense. You need ORGADMIN to manage multiple accounts.

upvoted 2 times

**mlaus** 3 years, 3 months ago

I think its BE: No, you need ACCOUNTADMIN role, or have the OWNERSHIP privilege. https://other-docs.snowflake.com/en/collaboration/provider-listings-creating-publishing.html#create-a-listing-for-the-sf-marketplace

upvoted 1 times

Question #155 *Topic 1*

When loading data into Snowflake via Snowpipe what is the compressed file size recommendation?

A. 10-50 MB

B. 100-250 MB

C. 300-500 MB

D. 1000-1500 MB

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

100 to 250 MB

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. 100-250 MB

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**EmiB** 3 years, 8 months ago

We recommend files at least above 10 MB on average, with files in the 100 to 250 MB range offering the best cost-to-performance ratio.
https://www.snowflake.com/blog/best-practices-for-data-ingestion

upvoted 1 times

**newbie2020** 3 years, 12 months ago

To optimize the number of parallel operations for a load, we recommend aiming to produce data files roughly 100MB to 250MB in size, compressed.
Splitting large files into a greater number of smaller files distributes the load among the servers in an active warehouse and increases performance.
https://community.snowflake.com/s/article/faq-does-snowflake-have-a-recommended-file-size-for-loading

upvoted 3 times

Question #156 *Topic 1*

Which Snowflake feature allows a user to substitute a randomly generated identifier for sensitive data, in order to prevent unauthorized users access to the data, before loading it into Snowflake?

A. External Tokenization

B. External Tables

C. Materialized Views

D. User-Defined Table Functions (UDTF)

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**_yyukta** Most Recent 2 years, 2 months ago

A. External Tokenization

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**Raju039** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/security-column-ext-token.html#external-tokenization

upvoted 2 times

**BigDataBB** 3 years, 3 months ago

https://docs.snowflake.com/en/user-guide/security-column-ext-token.html#external-tokenization

upvoted 3 times

Question #157 *Topic 1*

Which of the following are examples of operations that require a Virtual Warehouse to complete, assuming no queries have been executed previously? (Choose three.)

A. MIN( >)

B. COPY

C. SUM( >)

D. UPDATE

**Suggested Answer:** *B, C, D*

*Community vote distribution*

BCD (100%)

**Sk1020** Highly Voted 4 years ago

B,C,D are the answers. Min uses the Global cache / Service layer cache, hence it does not require the WH to execute query. So, A can be eliminated from answers.

upvoted 7 times

**SV1122** Highly Voted 3 years, 10 months ago

Was on exam Dec 31st 2022

upvoted 5 times

**SV1122** 3 years, 10 months ago

Was on exam Dec 31st 2022The min and max values of a colum are saved as metadata, so is not necessary use a WH to get their value

upvoted 5 times

**BigDataBB** Most Recent 4 months, 1 week ago

The min and max values of a colum are saved as metadata, so is not necessary use a WH to get their value

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B,C,D are options

upvoted 1 times

**_yyukta** 2 years, 8 months ago

BCD are correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**abbcia** 3 years, 7 months ago

min is in the cache

upvoted 1 times

**Jiaa** 3 years, 9 months ago

correct answer

upvoted 2 times

Question #158 *Topic 1*

What is the SNOWFLAKE.ACCOUNT_USAGE view that contains information about which objects were read by queries within the last 365 days (1 year)?

A. VIEWS_HISTORY

B. OBJECT_HISTORY

C. ACCESS_HISTORY

D. LOGIN_HISTORY

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

ACCESS_HISTORY

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C. ACCESS_HISTORY

upvoted 1 times

**Colourseun** 2 years, 11 months ago

The ANSWER is C. ACCESS_HISTORY. The ACCESS_HISTORY view contains information about which objects were accessed (read, written, or updated) by queries within the last 365 days (1 year). This includes information about the user who performed the access, the time of the access, and the object that was accessed.

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

Selected Answer: C

ideally Two options are correct Access_History and LOGIN_HISTORY. As it has been asked to select any one so going with C

https://docs.snowflake.com/en/sql-reference/account-usage.html#account-usage-views

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

ideally Two options are correct Access_History and LOGIN_HISTORY. As it has been asked to select any one so going with C

https://docs.snowflake.com/en/sql-reference/account-usage.html#account-usage-viewsSince asked "which objects were read by queries ", so the C must the only one answer.

upvoted 1 times

**otsumy** 3 years, 2 months ago

Since asked "which objects were read by queries ", so the C must the only one answer.

upvoted 1 times

**baig123** 3 years, 10 months ago

C is correct answer, there are no views with name views_history and object_history

upvoted 1 times

**kjdfsafc87456** 4 years, 1 month ago

Querying the ACCESS_HISTORY View This Account Usage view can be used to query the access history of Snowflake objects (e.g. table, view, column) within the last 365 days (1 year).

upvoted 4 times

Question #159 *Topic 1*

Which feature is only available in the Enterprise or higher editions of Snowflake?

A. Column-level security

B. SOC 2 type II certification

C. Multi-factor Authentication (MFA)

D. Object-level access control

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A.Column Level Security

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. Column-level security

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Sk1020** 4 years ago

Correct

upvoted 1 times

**greenlever** 4 years ago

correct

upvoted 1 times

Question #160 *Topic 1*

Will data cached in a warehouse be lost when the warehouse is resized?

A. Possibly, if the warehouse is resized to a smaller size and the cache no longer fits.

B. Yes, because the compute resource is replaced in its entirety with a new compute resource.

C. No, because the size of the cache is independent from the warehouse size.

D. Yes, because the new compute resource will no longer have access to the cache encryption key.

**Suggested Answer:** *A*

*Community vote distribution*

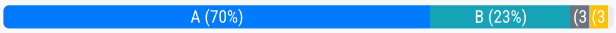


**William_20** Highly Voted 4 years, 2 months ago

Selected Answer: A

https://docs.snowflake.com/en/user-guide/warehouses-considerations.html#scaling-up-vs-scaling-out. At the "tip" it states: ... which CAN impact performance ... "Decreasing the size of a running warehouse removes compute resources from the warehouse. When the computer resources are removed, the cache associated with those resources is dropped, which can impact performance in the same way that suspending the warehouse can impact performance after it is resumed."

upvoted 19 times

**William_20** Highly Voted 4 years, 2 months ago

https://docs.snowflake.com/en/user-guide/warehouses-considerations.html#scaling-up-vs-scaling-out. At the "tip" it states: ... which CAN impact performance ... "Decreasing the size of a running warehouse removes compute resources from the warehouse. When the computer resources are removed, the cache associated with those resources is dropped, which can impact performance in the same way that suspending the warehouse can impact performance after it is resumed."Based on this explanation, answer is B. YES, cached data is dropped. right?Yes B is absolutely correctPossibly, if the warehouse is resized to a smaller size and the cache no longer fits.B is the right Answer

upvoted 19 times

**stornati** 3 years, 11 months ago

Based on this explanation, answer is B. YES, cached data is dropped. right?Yes B is absolutely correctPossibly, if the warehouse is resized to a smaller size and the cache no longer fits.

upvoted 4 times

**stornati** 3 years, 11 months ago

Based on this explanation, answer is B. YES, cached data is dropped. right?

upvoted 4 times

**d22770a** 2 years ago

Yes B is absolutely correct

upvoted 2 times

**BobFar** 3 years ago

Possibly, if the warehouse is resized to a smaller size and the cache no longer fits.

upvoted 1 times

**PrasantSadangi** 1 year, 7 months ago

B is the right Answer

upvoted 1 times

**Santa9os** Most Recent 1 year, 2 months ago

When you resize a warehouse (scale up or down), the cache is dropped because the warehouse is rebuilt with new compute resources.

upvoted 2 times

**d22770a** 2 years ago

https://docs.snowflake.com/en/user-guide/warehouses-considerations Decreasing the size of a running warehouse removes compute resources from the warehouse. When the computer resources are removed, the cache associated with those resources is dropped, which can impact performance in the same way that suspending the warehouse can impact performance after it is resumed.

upvoted 3 times

**Mallikharjuna452** 2 years, 3 months ago

A Is correct answer

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct Answer: A

upvoted 1 times

**sambathchandran** 2 years, 5 months ago

B. Yes, because the compute resource is replaced in its entirety with a new compute resource. When a Snowflake warehouse is resized, the current compute resources are replaced with new ones. This process clears the existing cache, and thus the cached data will be lost during the resizing operation.

upvoted 2 times

**Prashn** 2 years, 8 months ago

Decreasing the size of a running warehouse removes compute resources from the warehouse. When the computer resources are removed, the cache associated with those resources is dropped, which can impact performance in the same way that suspending the warehouse can impact performance after it is resumed. Keep this in mind when choosing whether to decrease the size of a running warehouse or keep it at the current size. In other words, there is a trade-off with regards to saving credits versus maintaining the cache.

upvoted 2 times

**_yyukta** 2 years, 8 months ago

A is correct

upvoted 1 times

**sandy91** 2 years, 9 months ago

Answer is A Decreasing the size of a running warehouse removes compute resources from the warehouse. When the computer resources are removed, the cache associated with those resources is dropped, which can impact performance in the same way that suspending the warehouse can impact performance after it is resumed.

upvoted 1 times

**Marge23** 3 years, 1 month ago

new compute resources will no longer have the data cache from old compute resources.

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

https://www.bing.com/search?q=Will+data+cached+in+a+warehouse+be+lost+when+the+warehouse+is+resized%3F%0D%0A&qs=n&form=QBRE&sp=-1&lq=1&pq=will+data+cached+in+a+warehouse+be+lost+when+the+warehouse+is+resized%3F&sc=1-71&sk=&cvid=EF9EDA8E014D4727A39525F970AE710F&ghsh=0&ghacc=0&ghpl=

upvoted 1 times

**Hopefan** 3 years, 5 months ago

my bad : Decreasing the size of a running warehouse removes compute resources from the warehouse. When the computer resources are removed, the cache associated with those resources is dropped, which can impact performance in the same way that suspending the warehouse can impact performance after it is resumed. so means possibly yes if it got reducerd so A is correct https://docs.snowflake.com/en/user-guide/warehouses-considerations

upvoted 3 times

**Hopefan** 3 years, 5 months ago

i believe C is correct @

upvoted 1 times

**learner2023** 3 years, 6 months ago

Decreasing the size of a running warehouse removes compute resources from the warehouse. When the computer resources are removed, the cache associated with those resources is dropped, which can impact performance in the same way that suspending the warehouse can impact performance after it is resumed. Keep this in mind when choosing whether to decrease the size of a running warehouse or keep it at the current size. In other words, there is a trade-off with regards to saving credits versus maintaining the cache.

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

https://docs.snowflake.com/en/user-guide/warehouses-considerations.html When the warehouse shrink it impacts

upvoted 1 times

**baig123** 3 years, 10 months ago

A is correct, only catch dropped if warehouse size is decreased.

upvoted 2 times

**db1290** 3 years, 11 months ago

possibly B

upvoted 3 times

Question #161 *Topic 1*

Which semi-structured file formats are supported when unloading data from a table? (Choose two.)

A. ORC

B. XML

C. Avro

D. Parquet

E. JSON

**Suggested Answer:** *D, E*

*Community vote distribution*

DE (100%)

**_yyukta** Most Recent 2 years, 2 months ago

D and E

upvoted 1 times

**kd_286** 2 years, 2 months ago

D:JSON

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**EmiB** 3 years, 1 month ago

D. Parquet E. JSON

upvoted 1 times

**SV1122** 3 years, 4 months ago

Structured: Delimited (CSV, TSV, etc.) Semi-structured: JSON, Parquet https://docs.snowflake.com/en/user-guide/data-unload-prepare.html

upvoted 2 times

Question #162 *Topic 1*

A running virtual warehouse is suspended.
What is the MINIMUM amount of time that the warehouse will incur charges for when it is restarted?

A. 1 second

B. 60 seconds

C. 5 minutes

D. 60 minutes

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B.60 sec

upvoted 1 times

**sambathchandran** 2 years, 5 months ago

In Snowflake, virtual warehouses are billed in one-minute increments. This means that when a suspended warehouse is restarted, it will incur charges for a minimum of 60 seconds, even if it is only used for a fraction of that time.

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: B

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. 60 seconds

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**MultiCloudIronMan** 4 years ago

Is thee right answer

upvoted 1 times

Question #163 *Topic 1*

What are the responsibilities of Snowflake's Cloud Service layer? (Choose three.)

A. Authentication

B. Resource management

C. Virtual warehouse caching

D. Query parsing and optimization

E. Query execution

F. Physical storage of micro-partitions

**Suggested Answer:** *A, B, D*

*Community vote distribution*

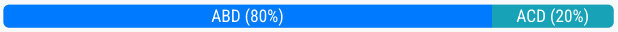


**pin2sadhu** Most Recent 1 year, 6 months ago

Correct

upvoted 1 times

**56d43d6** 1 year, 10 months ago

ABD. The virtual warehouse performs C virtual WH caching, not the cloud services layer. https://docs.snowflake.com/en/user-guide/intro-key-concepts#cloud-services

upvoted 1 times

**HICH9173** 2 years ago

acd https://docs.snowflake.com/en/user-guide/intro-key-concepts.html#cloud-services https://docs.snowflake.com/en/user-guide/cost-understanding-compute.html#cloud-service-credit-usage

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A,B,D ARE CORRECT

upvoted 1 times

**_yyukta** 2 years, 8 months ago

ABD are correct

upvoted 1 times

**happykeieiekek** 3 years, 3 months ago

ABD are correct

upvoted 1 times

**RajuNaik** 3 years, 6 months ago

ABD are correct options

upvoted 2 times

**KarBiswa** 3 years, 8 months ago

ABD are correct options

upvoted 3 times

**baig123** 3 years, 10 months ago

correct answer AB and D .

upvoted 3 times

**sdnv89** 3 years, 11 months ago

query execution happens in the warehouse

upvoted 2 times

**newbie2020** 3 years, 12 months ago

Sorry correcting, it should be ABD

upvoted 2 times

**newbie2020** 3 years, 12 months ago

Answer should be ACD. To support A&D : ---------- Services managed in this layer include: Authentication Infrastructure management Metadata management Query parsing and optimization Access control https://docs.snowflake.com/en/user-guide/intro-key-concepts.html#cloud-services To support C : ---------- This layer authenticates users, enforces security, performs query compilation and optimization, handles request query "caching", and more. https://docs.snowflake.com/en/user-guide/cost-understanding-compute.html#cloud-service-credit-usage

upvoted 4 times

**Sk1020** 4 years ago

Querying is the job of virtual warehouse

upvoted 2 times

**R12346** 4 years, 1 month ago

Should be ABD, but it is also used for querying metadata!

upvoted 2 times

**jveag** 4 years, 1 month ago

I thik same.

upvoted 1 times

**William_20** 4 years, 2 months ago

A, D for sure. B because of other options are not true.

upvoted 1 times

Question #164 *Topic 1*

How long is the Fail-safe period for temporary and transient tables?

A. There is no Fail-safe period for these tables.

B. 1 day

C. 7 days

D. 31 days

E. 90 days

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**pin2sadhu** Most Recent 1 year, 6 months ago

Correct , There is no fail safe period for Temporary and transient tables

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

There is no fail safe period for Temp and transient tables

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. There is no Fail-safe period for these tables.

upvoted 1 times

**singhks** 3 years, 2 months ago

'A' There is Zero (0) day Fail Safe period for temp and transient table

upvoted 1 times

**singhks** 3 years, 2 months ago

All answers are incorrect. There is Zero (0) day Fail Safe period for temp and transient tables.

upvoted 1 times

**bbrandhh** 3 years, 3 months ago

There was a similar question on the exam when I took it july 20th 2023

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**fahfouhi94** 3 years, 8 months ago

no fail safe for transient / temporary tables

upvoted 1 times

Question #165 *Topic 1*

Which command should be used to download files from a Snowflake stage to a local folder on a client's machine?

A. PUT

B. GET

C. COPY

D. SELECT

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B.GET COMMAND

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: B

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. GET

upvoted 1 times

**Colourseun** 2 years, 11 months ago

The command used to download files from a Snowflake stage to a local folder on a client's machine is GET. The GET command downloads the specified files from a Snowflake stage to the specified location on the client machine. You can use the GET command to download a single file, multiple files, or all of the files in a stage. GET @my_stage/myfile.csv file:///local/folder/myfile.csv;

upvoted 1 times

**bbrandhh** 3 years, 3 months ago

The correct answer is B for this question. An Important Note: Make sure you know how Snowflake delimits NULL and empty string values when downloading data from a stage. There was two questions on the exam when I took it asking that (july 2023) NULL - ,, Empty String - "

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**kjdfsafc87456** 4 years, 1 month ago

get is correct

upvoted 4 times

Question #166 *Topic 1*

How does Snowflake Fail-safe protect data in a permanent table?

A. Fail-safe makes data available up to 1 day, recoverable by user operations.

B. Fail-safe makes data available for 7 days, recoverable by user operations.

C. Fail-safe makes data available for 7 days, recoverable only by Snowflake Support.

D. Fail-safe makes data available up to 1 day, recoverable only by Snowflake Support.

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**_yyukta** Most Recent 2 years, 2 months ago

C. Fail-safe makes data available for 7 days, recoverable only by Snowflake Support.

upvoted 1 times

**Marge23** 2 years, 7 months ago

SF support only.

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**Sk1020** 3 years, 6 months ago

C is correct answer

upvoted 2 times

Question #167 *Topic 1*

A virtual warehouse is created using the following command:

Create warehouse my_WH with -
warehouse_size = MEDIUM
min_cluster_count = 1
max_cluster_count = 1
auto_suspend = 60
auto_resume = true;
The image below is a graphical representation of the warehouse utilization across two days.

What action should be taken to address this situation?

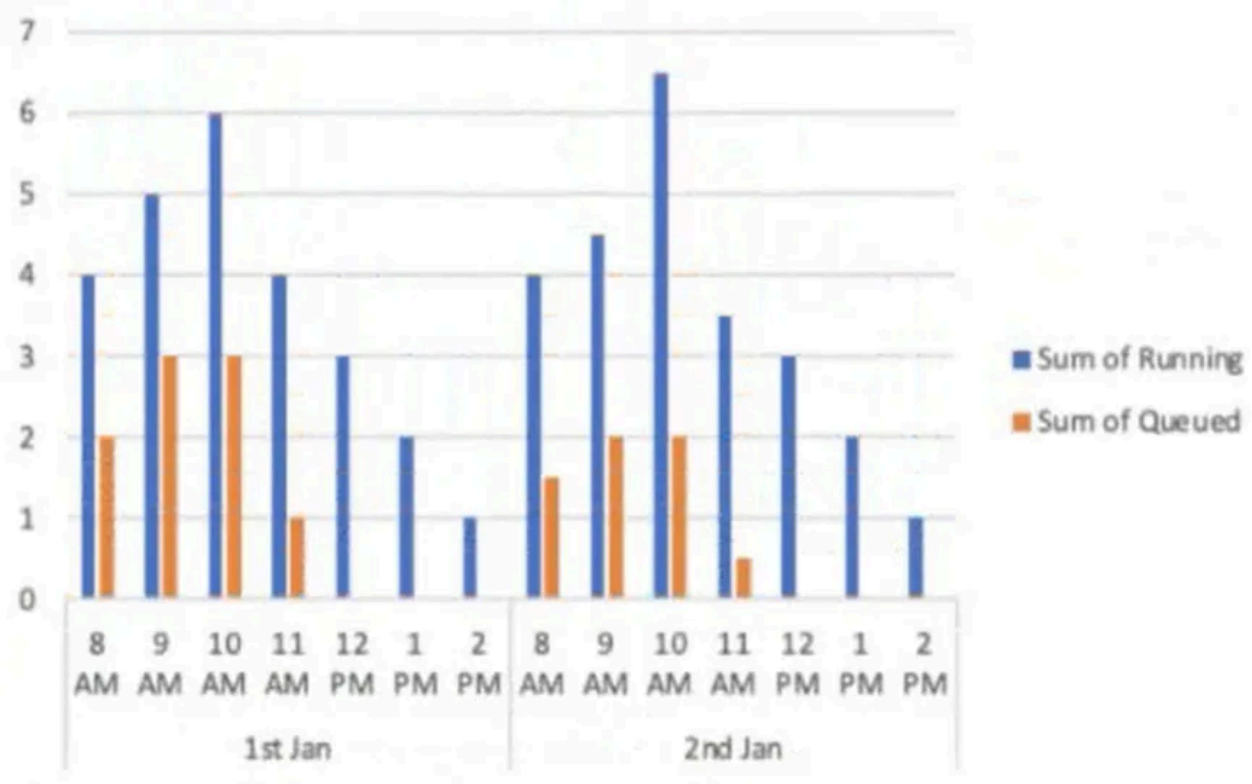


A. Increase the warehouse size from Medium to 2XL.

B. Increase the value for the parameter MAX_CONCURRENCY_LEVEL.

C. Configure the warehouse to a multi-cluster warehouse.

D. Lower the value of the parameter STATEMENT_QUEUED_TIMEOUT_IN_SECONDS.

**Suggested Answer:** *C*

*Community vote distribution*

C (90%) B (10%)

**NabilR** Highly Voted 3 years, 9 months ago

We can see that queries are queued starting from 4 queries running. Later 6 queries are running in concurrency. So queries are queued not for reaching the max possible number of queries, but for reaching the max Cluster utilization. Therefor, the best solution is to scale out (Answer C)

upvoted 5 times

**Shipra123** Most Recent 1 year, 11 months ago

Correct answer C.

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct Answer: B since it is already a multi cluster but having same min and max cluser size

upvoted 3 times

**pranalig** 2 years, 4 months ago

Correct Answer: B since it is already a multi cluster but having same min and max cluser sizeno, for WH to be multi cluster WH, max_cluster_count should be > 1 https://docs.snowflake.com/en/user-guide/warehouses-multiclusteryou are crazy

upvoted 3 times

**d22770a** 2 years ago

no, for WH to be multi cluster WH, max_cluster_count should be > 1 https://docs.snowflake.com/en/user-guide/warehouses-multicluster

upvoted 1 times

**sam_art** 2 years, 3 months ago

you are crazy

upvoted 4 times

**Chhatresh** 2 years, 6 months ago

If Queuing is the main concern here and if we set MAX_CONCURRENCY_LEVEL parameter to 8 or more, it can solve the Queuing problem but may lead to performance issue. queries will take longer time to complete as these share available compute resources. better solution is to scale out

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: B

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C. Configure the warehouse to a multi-cluster warehouse

upvoted 2 times

**Marge23** 3 years, 1 month ago

Selected Answer: B

Ans is B - MAX_CONCURRENCY_LEVEL The given is already a multicluster warehouse but only 1 cluster. so no need to configure the warehouse to a MC but to add cluster. Create warehouse my_WH with - warehouse_size = MEDIUM min_cluster_count = 1 max_cluster_count = 1 auto_suspend = 60 auto_resume = true;

upvoted 2 times

**Marge23** 3 years, 1 month ago

Ans is B - MAX_CONCURRENCY_LEVEL The given is already a multicluster warehouse but only 1 cluster. so no need to configure the warehouse to a MC but to add cluster. Create warehouse my_WH with - warehouse_size = MEDIUM min_cluster_count = 1 max_cluster_count = 1 auto_suspend = 60 auto_resume = true;ohh no! its not B its C but that is to add clusters.

upvoted 2 times

**Marge23** 3 years, 1 month ago

ohh no! its not B its C but that is to add clusters.

upvoted 2 times

**Def21** 3 years, 4 months ago

Nowadays default MAX_CONCURRENCY_LEVEL is 8 (previously 6?). Thus, concurrency is not an issue. We need multi-cluster https://docs.snowflake.com/en/sql-reference/parameters#max-concurrency-level

upvoted 4 times

**Leinho** 3 years, 5 months ago

I think is a very open question, what do you want to address? the performance or the queuing? If you want better performance in fewer queries so you should go for the MAX_CONCURRENCY_LEVEL, but if you want reduce the queuing you should look to the multi-clustering solution.Don't you think it is already a multi-cluster Maximized WH? It same value for min_cluster_count and max_cluster_countMy. bad max_cluster_count = 1, so it is NOT a muli-cluster WH, Ans is C

upvoted 2 times

**Leinho** 3 years, 5 months ago

I think is a very open question, what do you want to address? the performance or the queuing? If you want better performance in fewer queries so you should go for the MAX_CONCURRENCY_LEVEL, but if you want reduce the queuing you should look to the multi-clustering solution.

upvoted 2 times

**d22770a** 2 years ago

Don't you think it is already a multi-cluster Maximized WH? It same value for min_cluster_count and max_cluster_count

upvoted 1 times

**d22770a** 2 years ago

Don't you think it is already a multi-cluster Maximized WH? It same value for min_cluster_count and max_cluster_countMy. bad max_cluster_count = 1, so it is NOT a muli-cluster WH, Ans is C

upvoted 1 times

**d22770a** 2 years ago

My. bad max_cluster_count = 1, so it is NOT a muli-cluster WH, Ans is C

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

If auto scaling is enabled then it will be more flattened because edition is not mentioned

upvoted 1 times

**ajay_1233456** 3 years, 9 months ago

Better solution C

upvoted 1 times

**aalonso** 3 years, 10 months ago

Better solution C

upvoted 3 times

**baig123** 3 years, 10 months ago

correct answer C:

upvoted 1 times

**db1290** 3 years, 11 months ago

should be C

upvoted 1 times

**Sk1020** 4 years ago

Answer should be C as files can be seen in queue awaiting to be processed.

upvoted 1 times

**R12346** 4 years, 1 month ago

Should use multi-cluster and set the auto-scale parameter

upvoted 1 times

**jveag** 4 years, 1 month ago

Think same

upvoted 1 times

Question #168 *Topic 1*

Which minimum Snowflake edition allows for a dedicated metadata store?

A. Standard

B. Enterprise

C. Business Critical

D. Virtual Private Snowflake

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Virtual private snowflake

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. Virtual Private Snowflake

upvoted 1 times

**SV1122** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/intro-editions.html Dedicated metadata store and pool of compute resources (used in virtual warehouses) is offered through VPS

upvoted 3 times

**Sk1020** 4 years ago

explained in tabular format https://docs.snowflake.com/en/user-guide/intro-editions.html#enterprise-edition

upvoted 3 times

Question #169 *Topic 1*

Network policies can be set at which Snowflake levels? (Choose two.)

A. Role

B. Schema

C. User

D. Database

E. Account

F. Tables

**Suggested Answer:** *C, E*

*Community vote distribution*

CE (100%)

**kjdfsafc87456** Highly Voted 4 years, 1 month ago

Identifying a Network Policy Activated at the Account or User Level https://docs.snowflake.com/en/user-guide/network-policies.html#creating-network-policiesadd'l info.. activate n/w policy at 'security integration' level also (per documentation)

upvoted 9 times

**kjdfsafc87456** Highly Voted 4 years, 1 month ago

Selected Answer: CE

Identifying a Network Policy Activated at the Account or User Level https://docs.snowflake.com/en/user-guide/network-policies.html#creating-network-policies

upvoted 9 times

**subhataanu** 2 years, 5 months ago

add'l info.. activate n/w policy at 'security integration' level also (per documentation)

upvoted 1 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

C and E are Correct

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

Identifying a Network Policy Activated at the Account or User Level

upvoted 1 times

**Alex_Ovidiu** 2 years, 7 months ago

You are correct. Network policies can be set at two levels in Snowflake: D. Database E. Account Here's why the other options are not valid: A. Role: Network policies are not set at the role level. Roles are used to control what users can do within Snowflake, but they don't define network access restrictions. B. Schema: Network policies are not set at the schema level. Schemas are used to organize Snowflake objects, but they don't define network access restrictions. C. User: Network policies are not set at the user level. Users are assigned roles, and the roles determine the network access permissions. F. Tables: Network policies are not set at the table level. Tables store data, and network policies control access to Snowflake itself, not individual tables.

upvoted 1 times

**ASinanan** 2 years, 7 months ago

C and E Network policy precedence You can apply a network policy to an account, a security integration, or a user. If there are network policies applied to more than one of these, the most specific network policy overrides more general network policies. The following summarizes the order of precedence: Account Network policies applied to an account are the most general network policies. They are overridden by network policies applied to a security integration or user. Security Integration Network policies applied to a security integration override network policies applied to the account, but are overridden by a network policy applied to a user. User Network policies applied to a user are the most specific network policies. They override both accounts and security integrations.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C and E are correct

upvoted 1 times

**Colourseun** 2 years, 11 months ago

Network policies can be set at the Account and Role levels in Snowflake. So the answer is D. Database and E. Account. Network policies are used to control access to Snowflake resources from specific IP addresses or network ranges. You can create network policies at the account or role level. Account-level network policies apply to all users in the account, while role-level network policies apply only to users who are assigned that role.

upvoted 1 times

Question #170 *Topic 1*

What are the correct parameters for time travel and fail-safe in the Snowflake Enterprise Edition?

A. Default Time Travel Retention is set to 0 days. Maximum Time Travel Retention is 30 days. Fail Safe retention time is 1 day.

B. Default Time Travel Retention is set to 1 day. Maximum Time Travel Retention is 365 days. Fail Safe retention time is 7 days.

C. Default Time Travel Retention is set to 0 days. Maximum Time Travel Retention is 90 days. Fail Safe retention time is 7 days.

D. Default Time Travel Retention is set to 1 day. Maximum Time Travel Retention is 90 days. Fail Safe retention time is 7 days.

E. Default Time Travel Retention is set to 7 days. Maximum Time Travel Retention is 1 day. Fail Safe retention time is 90 days.

F. Default Time Travel Retention is set to 90 days. Maximum Time Travel Retention is 7 days. Fail Safe retention time is 356 days.

**Suggested Answer:** *D*

*Community vote distribution*




**Nagesh_CL** Highly Voted 4 years ago

Verified from snowflake documentation. D is the right answer.

upvoted 6 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

D is correct answer

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: D

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. Default Time Travel Retention is set to 1 day. Maximum Time Travel Retention is 90 days. Fail Safe retention time is 7 days.

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Def21** 3 years, 4 months ago

D still true.

upvoted 1 times

Question #171 *Topic 1*

Which of the following objects are contained within a schema? (Choose two.)

A. Role

B. Stream

C. Warehouse

D. External table

E. User

F. Share

**Suggested Answer:** *B, D*

*Community vote distribution*

BD (100%)

**pranalig** Most Recent 2 years, 4 months ago

Correct Answer: BD

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: BD

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B and D are correct

upvoted 1 times

**singhks** 3 years, 2 months ago

B and D are correct

upvoted 2 times

**alanhuangjob** 3 years, 7 months ago

Should be B & D

upvoted 1 times

**EmiB** 3 years, 8 months ago

Selected Answer: BD

BD is correct

upvoted 1 times

**EmiB** 3 years, 8 months ago

BD is correctA. Role (ACCOUNT) B. Stream (SCHEMA) C. Warehouse (ACCOUNT) D. External table (SCHEMA) E. User (ACCOUNT) F. Share (DATABASE)

upvoted 1 times

**EmiB** Highly Voted 3 years, 7 months ago

A. Role (ACCOUNT) B. Stream (SCHEMA) C. Warehouse (ACCOUNT) D. External table (SCHEMA) E. User (ACCOUNT) F. Share (DATABASE)

upvoted 23 times

**KarBiswa** 3 years, 8 months ago

external table and stream are the schema objects

upvoted 2 times

**ajay_1233456** 3 years, 9 months ago

BD is correct

upvoted 1 times

**seaun** 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-overview.html

upvoted 2 times

**greenlever** 4 years ago

BD is correct

upvoted 1 times

**R12346** 4 years, 1 month ago

not C at all

upvoted 1 times

**William_20** 4 years, 2 months ago

Certainly not C

upvoted 3 times

Question #172 *Topic 1*

Which of the following statements describe features of Snowflake data caching? (Choose two.)

A. When a virtual warehouse is suspended, the data cache is saved on the remote storage layer.

B. When the data cache is full, the least-recently used data will be cleared to make room.

C. A user can only access their own queries from the query result cache.

D. A user must set USE_METADATA_CACHE to TRUE to use the metadata cache in queries.

E. The RESULT_SCAN table function can access and filter the contents of the query result cache.

**Suggested Answer:** *B, E*

*Community vote distribution*




**DataEngineer7331** Highly Voted 3 years, 6 months ago

Option A is incorrect because when a virtual warehouse is suspended, the data cache is not saved on the remote storage layer. The data cache is cleared when a virtual warehouse is suspended and any data that needs to be cached is reloaded from the remote storage layer when the virtual warehouse is resumed. Option C is incorrect because the query result cache is a shared cache and all users can access the data that has been cached. There are no restrictions based on user access. Option D is incorrect because the metadata cache is used by default in queries and there is no need for a user to explicitly set USE_METADATA_CACHE to TRUE. Option B is correct because Snowflake automatically manages its data cache and evicts the least-recently used data when the cache becomes full. Option E is correct because the RESULT_SCAN table function can be used to query and filter the data that has been cached in the query result cache.Great explaination

upvoted 19 times

**DataEngineer7331** Highly Voted 3 years, 6 months ago

Selected Answer: BE

Option A is incorrect because when a virtual warehouse is suspended, the data cache is not saved on the remote storage layer. The data cache is cleared when a virtual warehouse is suspended and any data that needs to be cached is reloaded from the remote storage layer when the virtual warehouse is resumed.

Option C is incorrect because the query result cache is a shared cache and all users can access the data that has been cached. There are no restrictions based on user access.

Option D is incorrect because the metadata cache is used by default in queries and there is no need for a user to explicitly set USE_METADATA_CACHE to TRUE.

Option B is correct because Snowflake automatically manages its data cache and evicts the least-recently used data when the cache becomes full.

Option E is correct because the RESULT_SCAN table function can be used to query and filter the data that has been cached in the query result cache.

upvoted 18 times

**pranalig** 2 years, 4 months ago

Great explaination

upvoted 2 times

**William_20** 4 years, 2 months ago

result_scan: can return query results of past 24 hours: https://docs.snowflake.com/en/sql-reference/functions/result_scan.html

upvoted 10 times

**JRayan** Most Recent 1 year, 9 months ago

CE are the right option.

upvoted 1 times

**56d43d6** 1 year, 10 months ago

C is NOT correct: any user can retrieve results from the query result cache as long as the query is exactly the same as another user's query, and the underlying data has not changed. https://community.snowflake.com/s/article/Caching-in-the-Snowflake-Cloud-Data-Platform#:~:text=Snowflake%20Cache%20Layers&text=Result%20Cache:%20Which%20holds%20the,underlying%20data%20has%20not%20changed.

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct Answer: BE

upvoted 1 times

**sambathchandran** 2 years, 5 months ago

Snowflake's data cache follows a least-recently used (LRU) eviction policy. When the cache is full, the least-recently used data will be removed to make space for new data. The RESULT_SCAN table function allows users to access and filter the contents of the query result cache. This function is useful for examining and manipulating cached query results.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B and E are correct

upvoted 1 times

**Sk_3** 3 years ago

b,E RESULT_SCAN IS DIFFERENT FROM QUERY RESULT CACAHE

upvoted 1 times

**ajsharma22** 3 years, 2 months ago

"Only the user who ran the original query can use the RESULT_SCAN function to post-process the results of that query. Even another user with ACCOUNTADMIN privilege cannot access the results of another user's query by calling RESULT_SCAN." - https://docs.snowflake.com/en/sql-reference/functions/result_scan

upvoted 1 times

**dmitriypo** 3 years, 5 months ago

Only the user who ran the original query can use the RESULT_SCAN function to post-process the results of that query. Even another user with ACCOUNTADMIN privilege cannot access the results of another user's query by calling RESULT_SCAN. https://docs.snowflake.com/en/sql-reference/functions/result_scan#usage-notesYour statement is correct and that is why E is correct. but C is not correct. Query Result cache is different. This cached data is pulled when the same query is ran by any User/Role. but running queries specifically on this Query Result cache is done thru RESULT_SCAN. And this is only by the particular user who ran the initial query. this is because the user is now able to see the data.

upvoted 3 times

**dmitriypo** 3 years, 5 months ago

**Selected Answer: CE**

Only the user who ran the original query can use the RESULT_SCAN function to post-process the results of that query. Even another user with ACCOUNTADMIN privilege cannot access the results of another user's query by calling RESULT_SCAN. https://docs.snowflake.com/en/sql-reference/functions/result_scan#usage-notes

upvoted 3 times

**Rajivnb** 3 years ago

Your statement is correct and that is why E is correct. but C is not correct. Query Result cache is different. This cached data is pulled when the same query is ran by any User/Role. but running queries specifically on this Query Result cache is done thru RESULT_SCAN. And this is only by the particular user who ran the initial query. this is because the user is now able to see the data.

upvoted 1 times

**RajuNaik** 3 years, 6 months ago

CE correct answer

upvoted 1 times

**Kvk117** 3 years, 7 months ago

**Selected Answer: CE**

https://docs.snowflake.com/en/sql-reference/functions/result_scan

upvoted 2 times

**Kvk117** 3 years, 7 months ago

https://docs.snowflake.com/en/sql-reference/functions/result_scanThis is not right. The answer is B&E. The reasoning "Only the user who ran the original query can use the RESULT_SCAN function to post-process the results of that query." does not apply to using the cache in general. The role must have the same set of permissions to use the cache.

upvoted 2 times

**jawllalabad** 3 years, 6 months ago

This is not right. The answer is B&E. The reasoning "Only the user who ran the original query can use the RESULT_SCAN function to post-process the results of that query." does not apply to using the cache in general. The role must have the same set of permissions to use the cache.

upvoted 3 times

**brian_69420** 3 years, 7 months ago

CE. From SF docs: C: Only the user who ran the original query can use the RESULT_SCAN function to post-process the results of that query. Even another user with ACCOUNTADMIN privilege cannot access the results of another user's query by calling RESULT_SCAN. E: The query containing the RESULT_SCAN can include clauses, such as filters and ORDER BY clauses, that were not in the original query. This allows you to narrow down or modify the result set. https://docs.snowflake.com/en/sql-reference/functions/result_scan

upvoted 5 times

**baig123** 3 years, 10 months ago

B and E should be correct one.

upvoted 4 times

Question #173 *Topic 1*

A table needs to be loaded. The input data is in JSON format and is a concatenation of multiple JSON documents. The file size is 3 GB. A warehouse size S is being used. The following COPY INTO command was executed:
COPY INTO SAMPLE FROM @~/SAMPLE.JSON (TYPE=JSON)
The load failed with this error:
Max LOB size (16777216) exceeded, actual size of parsed column is 17894470.
How can this issue be resolved?

A. Compress the file and load the compressed file.

B. Split the file into multiple files in the recommended size range (100 MB - 250 MB).

C. Use a larger-sized warehouse.

D. Set STRIP_OUTER_ARRAY=TRUE in the COPY INTO command.

**Suggested Answer:** *D*

*Community vote distribution*

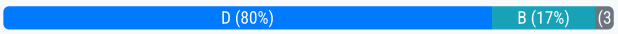


**0e504b5** Highly Voted 2 years, 9 months ago

D https://docs.snowflake.com/en/user-guide/semistructured-considerations#data-size-limitations Semi-structured Data Size Limitations A VARIANT can have a maximum size of up to 16 MB of uncompressed data. However, in practice, the maximum size is usually smaller due to internal overhead. The maximum size is also dependent on the object being stored. For more information, see VARIANT. In general, JSON data sets are a simple concatenation of multiple documents. The JSON output from some software is composed of a single huge array containing multiple records. There is no need to separate the documents with line breaks or commas, though both are supported. If the data exceeds 16 MB, enable the STRIP_OUTER_ARRAY file format option for the COPY INTO command to remove the outer array structure and load the records into separate table rowsNo, STRIP_OUTER_ARRAY is useful when loading a JSON file that contains a single outer array, but the question states that the file is a concatenation of multiple JSON documents. It does not address the LOB size error.

upvoted 7 times

**0e504b5** Highly Voted 2 years, 9 months ago

Selected Answer: D

D https://docs.snowflake.com/en/user-guide/semistructured-considerations#data-size-limitations

Semi-structured Data Size Limitations A VARIANT can have a maximum size of up to 16 MB of uncompressed data. However, in practice, the maximum size is usually smaller due to internal overhead. The maximum size is also dependent on the object being stored.

For more information, see VARIANT.

In general, JSON data sets are a simple concatenation of multiple documents. The JSON output from some software is composed of a single huge array containing multiple records. There is no need to separate the documents with line breaks or commas, though both are supported.

If the data exceeds 16 MB, enable the STRIP_OUTER_ARRAY file format option for the COPY INTO command to remove the outer array structure and load the records into separate table rows

upvoted 7 times

**BigDataBB** Most Recent 4 months, 1 week ago

No, STRIP_OUTER_ARRAY is useful when loading a JSON file that contains a single outer array, but the question states that the file is a concatenation of multiple JSON documents. It does not address the LOB size error.

upvoted 1 times

**alexhparreira** 5 months ago

Why Option B is correct: Snowflake enforces a strict maximum row size limit of 16 MB for semi-structured data types (such as VARIANT, OBJECT, or ARRAY). The error Max LOB size (16777216) exceeded means Snowflake attempted to read a single, unbroken JSON structure that was larger than 16 MB. Since the source data is a large 3 GB concatenation of multiple JSON documents, splitting this massive file into standard, smaller files (100 MB to 250 MB) allows Snowflake to correctly separate the distinct documents into manageable rows under the 16 MB limit during data ingestion.

upvoted 1 times

**niho1993** 1 year, 1 month ago

 Correct answer: B. Split the file into multiple files in the recommended size range.  D is only valid if your file has one large outer array, which your description ("concatenation of multiple JSON documents") says it does not.

upvoted 2 times

**PrasantSadangi** 1 year, 7 months ago

https://snowflake.discourse.group/t/what-json-format-does-strip_outer_array-support/7771

upvoted 1 times

**Siif97** 1 year, 9 months ago

The reason D. Set STRIP_OUTER_ARRAY=TRUE in the COPY INTO command is not the correct answer is because this option is used to handle JSON data that is wrapped in an outer array. It does not address the issue of exceeding the maximum LOB size. The error you encountered, "Max LOB size (16777216) exceeded," indicates that the size of a single JSON document within your file exceeds the maximum allowed size for a large object (LOB) in Snowflake So it's B

upvoted 2 times

**necroz** 1 year, 11 months ago

Removing brackets turns the JSON into an NDJSON, since every entry is valid json it can be parsed separately, the file won't be read in memory before being parsed

upvoted 1 times

**HICH9173** 2 years ago

it should be B split the files. https://community.snowflake.com/s/article/Max-LOB-size-exceeded#:~:text=When%20ingesting%20XML%20or%20JSON,of%20parsed%20column%20is%2017894470.

upvoted 2 times

**sam_art** 2 years, 3 months ago

Have you considered studying instead of guessing all the answers?

upvoted 2 times

**NeerajaVoleti** 2 years, 4 months ago

Correct Answer - B The error "Max LOB size (16777216) exceeded" indicates that the size of the parsed column exceeds the maximum size allowed for a single column value in Snowflake, which is 16 MB. To resolve this issue, the file should be split into multiple smaller files that are within the recommended size range of 100 MB to 250 MB. This will ensure that each JSON document within the files is smaller than the maximum LOB size allowed. Compressing the file, using a larger-sized warehouse, or setting STRIP_OUTER_ARRAY=TRUE will not resolve the issue of the column size exceeding the maximum allowed. References: COPY INTO Error during Structured Data Load: "Max LOB size (16777216) exceeded…"

upvoted 2 times

**pranalig** 2 years, 4 months ago

Correct Answer: D

upvoted 1 times

**GRKaushal** 2 years, 8 months ago

Ans A .....

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: D

upvoted 3 times

**_yyukta** 2 years, 8 months ago

D is correct

upvoted 3 times

**0e504b5** 2 years, 9 months ago

https://community.snowflake.com/s/article/Max-LOB-size-exceeded#:~:text=When%20ingesting%20XML%20or%20JSON,of%20parsed%20column%20is%2017894470. When ingesting XML or JSON documents, if the size of the compressed column data exceeds Snowflake's limit of 16MB [1], an error may occur: Max LOB size (16777216) exceeded, actual size of parsed column is 17894470.

upvoted 1 times

**oscarglob** 2 years, 10 months ago

The message also says that the column size that was read is a bit over 17MB. It's not the entire 3GB file. Wouldn't this create a problem by separating parts of the JSON that should be together?

upvoted 1 times

**Raju039** 3 years, 7 months ago

**Selected Answer: D**

D is correct.

upvoted 1 times

**Colourseun** 2 years, 11 months ago

The correct answer is B. Split the file into multiple files in the recommended size range (100 MB - 250 MB). The error message indicates that the size of the JSON data exceeds the maximum LOB size of 16 MB. This can happen if the JSON data contains large nested arrays or objects.

upvoted 1 times

**Arnab14** 3 years, 5 months ago

https://community.snowflake.com/s/article/Max-LOB-size-exceeded#:~:text=When%20ingesting%20XML%20or%20JSON,of%20parsed%20column%20is%2017894470.

upvoted 1 times

Question #174 *Topic 1*

What is a feature of a stored procedure in Snowflake?

A. They can be created as secure and hide the underlying metadata from all users.

B. They can access tables from a single database.

C. They can only contain a single SQL statement.

D. They can be created to run with a caller's rights or an owner's rights.

**Suggested Answer:** *D*

*Community vote distribution*

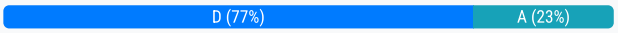


**PrashantGupta1616** 1 year, 11 months ago

answer is D

upvoted 1 times

**KarBiswa** Highly Voted 4 months ago

Selected Answer: B, D

The output of Flatten: SEQ A unique sequence number associated with the input record; the sequence is not guaranteed to be gap-free or ordered in any particular way. KEY For maps or objects, this column contains the key to the exploded value. PATH The path to the element within a data structure which needs to be flattened. INDEX The index of the element, if it is an array; otherwise NULL. VALUE The value of the element of the flattened array/object. THIS The element being flattened (useful in recursive flattening).

upvoted 18 times

**_yyukta** 2 years, 2 months ago

A. They can be created as secure and hide the underlying metadata from all users.

upvoted 1 times

**SV1122** 4 months ago

Was on exam Dec 31st 2022

upvoted 3 times

**Colourseun** 2 years, 5 months ago

The answer is D. They can be created to run with a caller's rights or an owner's rights. Stored procedures in Snowflake can be created to run with either the caller's rights or the owner's rights. This allows you to control the level of access that users have to the data and resources that are used in the stored procedure.

upvoted 1 times

**AlexbDku** 4 months ago

Selected Answer: B, D

https://docs.snowflake.com/en/sql-reference/functions/flatten.html

upvoted 3 times

**BobFar** 2 years, 6 months ago

Answer is D , I am not sure about A. For example, for a secure function or procedure, information omitted for unauthorized users includes its: Body (the handler code that comprises its logic) List of imports Handler name Packages list Unauthorized users will still be able to see information that includes its: Parameter types Return type Handler language Null handling Volatility https://docs.snowflake.com/en/developer-guide/secure-udf-procedure#label-secure-limiting-visibility-definition But D is definitely true A stored procedure runs with either the caller's rights or the owner's rights. It cannot run with both at the same time. This topic describes the differences between a caller's rights stored procedure and an owner's rights stored procedure. https://docs.snowflake.com/en/developer-guide/stored-procedure/stored-procedures-rights

upvoted 1 times

**tejarun** 4 months ago

Selected Answer: B, D

https://docs.snowflake.com/en/sql-reference/functions/flatten.html

upvoted 3 times

**Marge23** 2 years, 7 months ago

A - https://docs.snowflake.com/en/developer-guide/stored-procedure/stored-procedures-overview

upvoted 1 times

**SoftwareEngineerGoneData** Most Recent 4 months ago

Selected Answer: B, D

https://docs.snowflake.com/en/sql-reference/functions/flatten#output

upvoted 2 times

**singhks** 2 years, 8 months ago

A and D https://docs.snowflake.com/developer-guide/stored-procedure/stored-procedures-overview and https://docs.snowflake.com/developer-guide/secure-udf-procedure

upvoted 1 times

**SoftwareEngineerGoneData** 2 years, 8 months ago

https://docs.snowflake.com/en/sql-reference/stored-procedures-overview#what-is-a-stored-procedure

upvoted 1 times

**malexandersalazar** 2 years, 11 months ago

I think that the "Correct Answers" are designed to fail hahaha. D is the correct.

upvoted 2 times

**Raju039** 3 years, 1 month ago

D is correct...https://docs.snowflake.com/en/sql-reference/stored-procedures-rights

upvoted 1 times

**EmiB** 3 years, 2 months ago

A - False - When you specify that the UDF or procedure is secure, these details are visible only to authorized users – in other words, to users who are granted a role that owns the function.

upvoted 1 times

**KarBiswa** 3 years, 2 months ago

Invokers and caller rights

upvoted 1 times

**SV1122** 3 years, 4 months ago

Stored procs can be created to be secure https://docs.snowflake.com/en/developer-guide/secure-udf-procedure.html#determining-if-a-udf-or-procedure-is-securebut can't hide the underlying metadata from all users. Here "all" users is the catch. The role that created can view the definition of stored procAccording to documentation is D.

upvoted 1 times

**SV1122** 3 years, 4 months ago

but can't hide the underlying metadata from all users. Here "all" users is the catch. The role that created can view the definition of stored procAccording to documentation is D.

upvoted 2 times

**sakis213** 3 years, 4 months ago

According to documentation is D.

upvoted 1 times

**baig123** 3 years, 4 months ago

should be D

upvoted 1 times

**db1290** 3 years, 5 months ago

Should be D

upvoted 1 times

**Sk1020** 3 years, 6 months ago

Caller's Owner's is correct . Secured option is good for views, Materialized views.

upvoted 1 times

**William_20** 3 years, 8 months ago

https://docs.snowflake.com/en/sql-reference/stored-procedures-rights.html

upvoted 3 times

Question #175 *Topic 1*

Which columns are part of the result set of the Snowflake LATERAL FLATTEN command? (Choose two.)

A. CONTENT

B. PATH

C. BYTE_SIZE

D. INDEX

E. DATATYPE

**Suggested Answer:** *B, D*

*Community vote distribution*

BD (100%)

**KarBiswa** Highly Voted 3 years, 8 months ago

The output of Flatten: SEQ A unique sequence number associated with the input record; the sequence is not guaranteed to be gap-free or ordered in any particular way. KEY For maps or objects, this column contains the key to the exploded value. PATH The path to the element within a data structure which needs to be flattened. INDEX The index of the element, if it is an array; otherwise NULL. VALUE The value of the element of the flattened array/object. THIS The element being flattened (useful in recursive flattening).

upvoted 18 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

PATH,INDEX

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B and D are correct

upvoted 1 times

**SoftwareEngineerGoneData** 3 years, 2 months ago

https://docs.snowflake.com/en/sql-reference/functions/flatten#output

upvoted 2 times

**VEEDOL** 3 years, 5 months ago

answer - BD In Snowflake, the LATERAL FLATTEN command is used to unnest semi-structured data stored in variant or object columns. When you execute the LATERAL FLATTEN command, it generates a new virtual table with the unnested data. The result set of the LATERAL FLATTEN command includes the following columns: KEY: This column contains the key or path to the unnested element within the original semi-structured data. INDEX: This column represents the index of the unnested element within the array or object. ELEMENT: This column contains the value of the unnested element. These three columns make up the result set of the LATERAL FLATTEN command and provide the necessary information to access and process the unnested data.

upvoted 2 times

**akhandelwal680** 3 years, 5 months ago

A. CONTENT and B. PATH. The LATERAL FLATTEN command expands a VARIANT, OBJECT, or ARRAY column into multiple rows, where CONTENT represents the expanded values and PATH represents the path to the expanded values within the original column. The columns C. BYTE_SIZE, D. INDEX, and E. DATATYPE are not part of the result set of the LATERAL FLATTEN command.

upvoted 1 times

**Sharan009** 3 years, 9 months ago

The result value of flatten command is SEQ, PATH, KEY, INDEX, VALUE, THIS

upvoted 2 times

**SV1122** 3 years, 10 months ago

Was on exam Dec 31st 2022

upvoted 3 times

**AlexbDku** 3 years, 10 months ago

https://docs.snowflake.com/en/sql-reference/functions/flatten.html

upvoted 3 times

**halol** 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/functions/flatten.html

upvoted 2 times

**tejarun** 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/functions/flatten.html

upvoted 3 times

Question #176 *Topic 1*

What is the minimum Snowflake edition required to create a materialized view?

A. Standard Edition

B. Enterprise Edition

C. Business Critical Edition

D. Virtual Private Snowflake Edition

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Enterprise edition

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

The minimum Snowflake edition required to create a materialized view is: B. Enterprise Edition Materialized views are supported starting from the Enterprise Edition of Snowflake.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Enterprise Edition

upvoted 1 times

**Marge23** 3 years, 1 month ago

enterprise B

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**SV1122** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/intro-editions.html

upvoted 2 times

**halol** 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/sql/create-materialized-view.html#:~:text=Materialized%20views%20require%20Enterprise%20Edition,upgrading%2C%20please%20contact%20Snowflake%20Support.

upvoted 1 times

Question #177 *Topic 1*

Which Snowflake function will interpret an input string as a JSON document, and produce a VARIANT value?

A. parse_json()

B. json_extract_path_text()

C. object_construct()

D. flatten

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**tejarun** Highly Voted 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/functions/parse_json.html

upvoted 8 times

**EmiB** Highly Voted 4 months ago

Selected Answer: A

A - correct B - serverless has no minimum (virtual warehouse has 60 sec) C - parameter doesn't exists D - not 10% of MONTHLY, but DAILY usage - Usage for cloud services is charged only if the daily consumption of cloud services exceeds 10% of the daily usage of virtual warehouses.

upvoted 13 times

**Mallikharjuna452** 2 years, 3 months ago

parse_json()

upvoted 1 times

**muks86** 4 months ago

The most confusing question

upvoted 3 times

**pranalig** 2 years, 4 months ago

Correct Answer: A

upvoted 1 times

**n21007** 4 months ago

Charges for serverless features are calculated based on total usage of snowflake-managed compute resources measured in compute-hours. Compute-Hours are calculated on a per second basis, rounded up to the nearest whole second. The number of credits consumed per compute hour varies depending on the serverless feature. To learn how many credits are consumed by a serverless feature, refer to the “Serverless Feature Credit Table” in the Snowflake service consumption table. https://docs.snowflake.com/en/user-guide/cost-understanding-compute#serverless-credit-usage

upvoted 3 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: B

upvoted 1 times

**Geera** Most Recent 4 months ago

I wish they weren't billed...

upvoted 3 times

**_yyukta** 2 years, 8 months ago

A parse_json()

upvoted 1 times

**YokoSumiGaeshi** 4 months ago

Selected Answer: A, I

Even though D is true, it doesn't answer the question, so I'd say A but it's not really clear.

upvoted 2 times

**Marge23** 3 years, 1 month ago

correct

upvoted 1 times

**singhks** 3 years, 2 months ago

A is the answer. https://docs.snowflake.com/en/sql-reference/functions/parse_json

upvoted 1 times

**examtopics_strata** 3 years, 3 months ago

A is correct

upvoted 1 times

**qaiserlatif** 3 years, 8 months ago

Interprets an input string as a JSON document, producing a VARIANT value. Syntax PARSE_JSON( )

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

parse_json

upvoted 1 times

**ViniJsr** 3 years, 10 months ago

The correct answer is A.

upvoted 1 times

**SV1122** 3 years, 10 months ago

Was on exam Dec 31st 2022

upvoted 2 times

**paknadeem** 3 years, 11 months ago

The correct answer is A

upvoted 3 times

**KeshavaMugulur** 3 years, 11 months ago

This should A for sure

upvoted 2 times

Question #178 *Topic 1*

How are serverless features billed?

A. Per second multiplied by an automatic sizing for the job

B. Per minute multiplied by an automatic sizing for the job, with a minimum of one minute

C. Per second multiplied by the size, as determined by the SERVERLESS_FEATURES_SIZE account parameter

D. Serverless features are not billed, unless the total cost for the month exceeds 10% of the warehouse credits, on the account

**Suggested Answer:** *A*

*Community vote distribution*

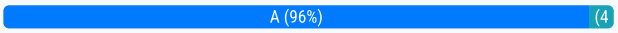


**EmiB** Highly Voted 3 years, 7 months ago

A - correct B - serverless has no minimum (virtual warehouse has 60 sec) C - parameter doesn't exists D - not 10% of MONTHLY, but DAILY usage - Usage for cloud services is charged only if the daily consumption of cloud services exceeds 10% of the daily usage of virtual warehouses.

upvoted 14 times

**hotflake** Most Recent 2 years, 3 months ago

Billed per second, per core

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

Charges for serverless features are calculated based on total usage of snowflake-managed compute resources measured in compute-hours. Compute-Hours are calculated on a per second basis, rounded up to the nearest whole second. The number of credits consumed per compute hour varies depending on the serverless feature.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. Per second multiplied by an automatic sizing for the job

upvoted 1 times

**rooo4567** 3 years, 7 months ago

D for fure

upvoted 1 times

**bibisa** 3 years, 8 months ago

Charges for serverless features are calculated based on total usage of snowflake-managed compute resources measured in compute-hours. Compute-Hours are calculated on a per second basis, rounded up to the nearest whole second. The number of credits consumed per compute hour varies depending on the serverless feature.

upvoted 1 times

**mighty_kafta** 3 years, 8 months ago

Very confusing question but i think A is the answer

upvoted 1 times

**ShagunMittal** 3 years, 8 months ago

Serverless features are billed per second. And there is no parameter by name SERVERLESS_FEATURES_SIZE. So A should be correct.

upvoted 1 times

**muks86** 3 years, 8 months ago

The most confusing question

upvoted 3 times

**n21007** 3 years, 8 months ago

Charges for serverless features are calculated based on total usage of snowflake-managed compute resources measured in compute-hours. Compute-Hours are calculated on a per second basis, rounded up to the nearest whole second. The number of credits consumed per compute hour varies depending on the serverless feature. To learn how many credits are consumed by a serverless feature, refer to the “Serverless Feature Credit Table” in the Snowflake service consumption table. https://docs.snowflake.com/en/user-guide/cost-understanding-compute#serverless-credit-usage

upvoted 3 times

**YokoSumiGaeshi** 3 years, 8 months ago

Even though D is true, it doesn't answer the question, so I'd say A but it's not really clear.

upvoted 2 times

**KarBiswa** 3 years, 8 months ago

its the answer

upvoted 1 times

**JRayan** 3 years, 8 months ago

Serverless compute does not factor into the 10% adjustment for cloud services.https://docs.snowflake.com/en/user-guide/cost-understanding-compute.htmlIt should be C and definitely not D

upvoted 1 times

**JRayan** 3 years, 8 months ago

Serverless compute does not factor into the 10% adjustment for cloud services.

upvoted 1 times

**JRayan** 3 years, 8 months ago

https://docs.snowflake.com/en/user-guide/cost-understanding-compute.html

upvoted 1 times

**JRayan** 3 years, 8 months ago

It should be C and definitely not D

upvoted 1 times

**SalihAlic** 3 years, 9 months ago

I think A It says that it is billed based on compute-hours which are charged per second, with Snowflake determining the size which depends on the compute needs.

upvoted 2 times

**Jiaa** 3 years, 9 months ago

I feel it's more like C https://docs.snowflake.com/en/user-guide/cost-understanding-compute.html#serverless-credit-usage

upvoted 1 times

**Geera** 3 years, 9 months ago

I wish they weren't billed...

upvoted 3 times

**halol** 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/cost-understanding-compute.html#:~:text=for%20Cloud%20Services-,Usage,-Usage%20for%20cloud should be D but if the cost daily exceeds 10%

upvoted 2 times

Question #179 *Topic 1*

Which Snowflake architectural layer is responsible for a query execution plan?

A. Compute

B. Data storage

C. Cloud services

D. Cloud provider

**Suggested Answer:** *C*

*Community vote distribution*

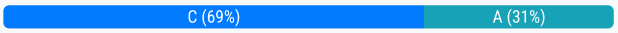


**Santa9os** Most Recent 1 year, 2 months ago

Cloud Services layer: Handles authentication, metadata management, query parsing, and query optimization (which includes building the query execution plan).

upvoted 1 times

**56d43d6** 1 year, 10 months ago

Query execution plan is executed by the cloud services layer. https://docs.snowflake.com/en/sql-reference/sql/explain

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

C.Cloud service

upvoted 1 times

**Lematthew31** 2 years, 4 months ago

This is A not D, because it's the compute layer that execute, cloud services will be more for query parsing

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

The Snowflake architectural layer responsible for a query execution plan is: C. Cloud services The Cloud Services layer in Snowflake handles activities such as query optimization, compilation, and execution planning.

upvoted 1 times

**nexerSnow** 2 years, 5 months ago

Query execution is performed in the processing layer. Snowflake processes queries using “virtual warehouses”. Each virtual warehouse is an MPP compute cluster composed of multiple compute nodes allocated by Snowflake from a cloud provider.

upvoted 1 times

**ede3703** 2 years, 6 months ago

Compute Layer

upvoted 1 times

**sushant02** 2 years, 6 months ago

why not A? Compute layer

upvoted 3 times

**_yyukta** 2 years, 8 months ago

C. Cloud services

upvoted 1 times

**4db5e62** 2 years, 9 months ago

Made a mistake buying this set, all answers are wrong. Will never ever recommend this to anyone. Refrain from using this.

upvoted 4 times

**0e504b5** 2 years, 9 months ago

C https://docs.snowflake.com/en/user-guide/cost-understanding-compute#label-cloud-services-credit-usage https://docs.snowflake.com/en/user-guide/intro-key-concepts

upvoted 1 times

**4db5e62** 2 years, 9 months ago

Made a mistake buying this set, all answers are wrong. Will never ever recommend this to anyone. Refrain from using this.At least the discussion section under each question looks very valuable from learning perspective. You have a chance to see other opinions, sometimes different than yours which urges you to your own deeper research and better understanding.

upvoted 4 times

**rogerek077** 2 years, 7 months ago

At least the discussion section under each question looks very valuable from learning perspective. You have a chance to see other opinions, sometimes different than yours which urges you to your own deeper research and better understanding.

upvoted 4 times

**Matsalearnings** 3 years, 1 month ago

why not a Cloud Provider as per the below- Query execution is performed in the processing layer. Snowflake processes queries using “virtual warehouses”. Each virtual warehouse is an MPP compute cluster composed of multiple compute nodes allocated by Snowflake from a cloud provider.

upvoted 1 times

**Thaiya** 3 years, 1 month ago

cloud services, Query Processing and Data storage are the layers. Must be A compute as its not any of the other 3. It must imply query processing layer.

upvoted 1 times

**Laliye** 3 years, 7 months ago

cloud services layer. C is the answer

upvoted 3 times

**KarBiswa** 3 years, 8 months ago

Cloud services layer

upvoted 1 times

**sakis213** 3 years, 10 months ago

Query execution plan, does not need an active virtual warehouse to complete as it only compiles the SQL queries and consumes only Cloud service credits, since its a metadata operation. I would go with C.

upvoted 2 times

**Naah** 3 years, 10 months ago

Cloud Services The cloud services layer is a collection of services that coordinate activities across Snowflake. These services tie together all of the different components of Snowflake in order to process user requests, from login to query dispatch. The cloud services layer also runs on compute instances provisioned by Snowflake from the cloud provider.

Services managed in this layer include:

Authentication

Infrastructure management

Metadata management

Query parsing and optimization

Access control

upvoted 4 times

**chandra535** 3 years, 10 months ago

correct answer is C

upvoted 1 times

Question #180 *Topic 1*

When unloading to a stage, which of the following is a recommended practice or approach?

A. Set SINGLE = TRUE for larger files.

B. Use OBJECT_CONSTRUCT(*) when using Parquet.

C. Avoid the use of the CAST function.

D. Define an individual file format.

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**SV1122** Highly Voted 3 years, 10 months ago

Data Unloading Considerations: A. Defining a File Format: File format defines the type of data to be unloaded into the stage or S3. It is best practice to define an individual file format when regularly used to unload a certain type of data based on the characteristics of the file needed.
https://community.snowflake.com/s/article/Best-Practices-for-Data-Unloading

upvoted 10 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Define an Individual format

upvoted 1 times

**pranalig** 2 years, 4 months ago

correct answer:D

upvoted 1 times

**Chhatresh** 2 years, 6 months ago

why not single=true for larger files?

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: D

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. Define an individual file format.

upvoted 1 times

**Sk_3** 3 years ago

y this is wrong Use OBJECT_CONSTRUCT(*) when using Parquet...?

upvoted 2 times

**psilvabr** 3 years, 8 months ago

D is the anwser

upvoted 2 times

**KarBiswa** 3 years, 8 months ago

D is best

upvoted 1 times

**NabilR** 3 years, 9 months ago

define an individual file format is a best practice

upvoted 1 times

**Lavpak** 3 years, 11 months ago

Avoid use of the CAST function is not an answer https://docs.snowflake.com/en/user-guide/data-unload-considerations.html

upvoted 3 times

Question #181 *Topic 1*

Which SQL commands, when committed, will consume a stream and advance the stream offset? (Choose two.)

A. UPDATE TABLE FROM STREAM

B. SELECT FROM STREAM

C. INSERT INTO TABLE SELECT FROM STREAM

D. ALTER TABLE AS SELECT FROM STREAM

E. BEGIN COMMIT

**Suggested Answer:** *A, C*

*Community vote distribution*

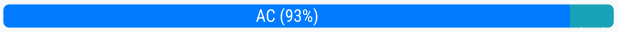


**Mallikharjuna452** Most Recent 2 years, 3 months ago

A and C are correct

upvoted 1 times

**pranalig** 2 years, 4 months ago

correct answer:AC

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A and C are correct

upvoted 1 times

**Colourseun** 2 years, 11 months ago

The following SQL commands, when committed, will consume a stream and advance the stream offset: A: SELECT FROM STREAM B: UPDATE TABLE FROM STREAM

upvoted 1 times

**ludaka** 3 years, 1 month ago

A and C

upvoted 2 times

**SushY18** 3 years, 6 months ago

https://docs.snowflake.com/en/user-guide/streams-intro

To advance the offset of a stream to the current table version without consuming the change data in a DML operation, complete either of the following actions:

Recreate the stream (using the CREATE OR REPLACE STREAM syntax).

Insert the current change data into a temporary table. In the INSERT statement, query the stream but include a WHERE clause that filters out all of the change data (e.g. WHERE 0 = 1).

upvoted 1 times

**SushY18** 3 years, 6 months ago

https://docs.snowflake.com/en/user-guide/streams-intro To advance the offset of a stream to the current table version without consuming the change data in a DML operation, complete either of the following actions: Recreate the stream (using the CREATE OR REPLACE STREAM syntax). Insert the current change data into a temporary table. In the INSERT statement, query the stream but include a WHERE clause that filters out all of the change data (e.g. WHERE 0 = 1).choose CD

upvoted 1 times

**SushY18** 3 years, 6 months ago

choose CD

upvoted 2 times

**Raju039** 3 years, 7 months ago

AC The stream position (i.e. offset) is advanced when the stream is used in a DML statement. The position is updated at the end of the transaction to the beginning timestamp of the transaction. The stream describes change records starting from the current position of the stream and ending at the current transactional timestamp. To ensure multiple statements access the same change records in the stream, surround them with an explicit transaction statement (BEGIN .. COMMIT). An explicit transaction locks the stream, so that DML updates to the source object are not reported to the stream until the transaction is committed.

upvoted 4 times

**EmiB** 3 years, 7 months ago

Could be ACE: A, C - The stream position (i.e. offset) is advanced when the stream is used in a DML statement. E - To ensure multiple statements access the same change records in the stream, surround them with an explicit transaction statement (BEGIN .. COMMIT).

upvoted 1 times

**psilvabr** 3 years, 8 months ago

DMLs advance Stream, so it is A and C.

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

If committed after update it will move the offset position and same for insert

upvoted 1 times

**mlaus** 3 years, 9 months ago

Insert, Delete, Update, so its AC

upvoted 1 times

**Mubin_Sikilkar** 3 years, 9 months ago

The stream position is advanced in DML statements i.e (Update,Insert) Hence A&C is the correct answer

upvoted 4 times

**SV1122** 3 years, 10 months ago

The stream position (i.e. offset) is advanced when the stream is used in a DML statement. Insert, Update

upvoted 2 times

**paknadeem** 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/sql/create-stream.html

upvoted 3 times

**halol** 3 years, 11 months ago

A&C DML

upvoted 4 times

Question #182 *Topic 1*

Which methods can be used to delete staged files from a Snowflake stage? (Choose two.)

A. Use the DROP [file] command after the load completes.

B. Specify the TEMPORARY option when creating the file format.

C. Specify the PURGE copy option in the COPY INTO [table] command.

D. Use the REMOVE command after the load completes.

E. Use the DELETE LOAD HISTORY command after the load completes.

**Suggested Answer:** *C, D*

*Community vote distribution*




**GiselaS** Highly Voted 3 years, 7 months ago
Purge in COPY INTO command or REMOVE after that process
upvoted 5 times

**PrasantSadangi** Most Recent 1 year, 7 months ago
Purge and remove
upvoted 1 times

**a8c5bbe** 1 year, 11 months ago
purge and remove
upvoted 1 times

**sanjithraju39** 1 year, 11 months ago
cd is the answer
upvoted 1 times

**pranalig** 2 years, 4 months ago
CD correct
upvoted 1 times

**82c6010** 2 years, 5 months ago
C and D
upvoted 1 times

**gunasekarantjk1996** 2 years, 5 months ago
CD correct
upvoted 1 times

**_yyukta** 2 years, 8 months ago
C and D are correct
upvoted 1 times

**Marge23** 3 years, 1 month ago
A is to delete a stage by using command DROP @mystage;
upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago
correct
upvoted 2 times

**weuz** 3 years, 7 months ago
PURGE and REMOVE are the correct answer
upvoted 1 times

**c0d3g** 3 years, 7 months ago
PURGE and REMOVE
upvoted 1 times

**psilvabr** 3 years, 8 months ago

CD correct answer.

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

CD correct anser for C - https://docs.snowflake.com/en/sql-reference/sql/copy-into-table.html

upvoted 1 times

**baig123** 3 years, 10 months ago

C and D should be correct answer.

upvoted 3 times

**halol** 3 years, 11 months ago

d: https://docs.snowflake.com/en/user-guide/data-load-internal-tutorial-remove-data-files.html c: https://docs.snowflake.com/en/user-guide/data-load-local-file-system-copy.html#:~:text=Staged%20files%20can%20be%20deleted%20from%20a%20Snowflake,command%20to%20remove%20the%20files%20in%20the%20stage.

upvoted 4 times

**db1290** 3 years, 11 months ago

C and D

upvoted 3 times

Question #183 *Topic 1*

On which of the following cloud platforms can a Snowflake account be hosted? (Choose three.)

A. Amazon Web Services

B. Private Virtual Cloud

C. Oracle Cloud

D. Microsoft Azure Cloud

E. Google Cloud Platform

F. Alibaba Cloud

**Suggested Answer:** *A, D, E*

*Community vote distribution*

ADE (86%) ACE (14%)

**a8c5bbe** Most Recent 1 year, 11 months ago

A,D and E are correct

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A,D,E are correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

ADE are correct

upvoted 1 times

**Marge23** 3 years, 1 month ago

AWS, GCP, AZURE

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**psilvabr** 3 years, 8 months ago

ADE is the way

upvoted 1 times

**halol** 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/intro-cloud-platforms.html#:~:text=A%20Snowflake%20account%20can%20be,Microsoft%20Azure%20(Azure)

upvoted 2 times

Question #184 *Topic 1*

What Snowflake role must be granted for a user to create and manage accounts?

A. ACCOUNTADMIN

B. ORGADMIN

C. SECURITYADMIN

D. SYSADMIN

**Suggested Answer:** *B*

*Community vote distribution*

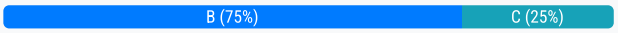


**SV1122** Highly Voted 3 years, 10 months ago

Was on exam Dec 31st 2022

upvoted 6 times

**56d43d6** Most Recent 1 year, 10 months ago

B. An "account" is not a user. An account is below Organization and above Database in the hierarchy. https://docs.snowflake.com/en/user-guide/organizations-manage-accounts

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

ORGADMIN IS CORRECT

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. ORGADMIN

upvoted 1 times

**singhks** 3 years, 2 months ago

The key phrase is : "a user to create and manage accounts? An ORGADMIN can create / manage Accounts and hence AccountADMINs.

upvoted 2 times

**lukalooko** 3 years, 3 months ago

The question is for 'account' , not creating a role or user. An 'account' can be thought of as the url you are using to login to your database. https://docs.snowflake.com/en/user-guide/organizations-manage-accounts

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

https://www.bing.com/search?q=What+Snowflake+role+must+be+granted+for+a+user+to+create+and+manage+accounts%3F&qs=n&form=QBRE&sp=-1&lq=1&pq=what+snowflake+role+must+be+granted+for+a+user+to+create+and+manage+accounts%3F&sc=1-77&sk=&cvid=A5167129668F46538ECF45FD52DDE182&ghsh=0&ghacc=0&ghpl=

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Selected Answer: B

Correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

CorrectI changed my mind it C, SecurityAdmin that inherit the priveleges of UserAdmin

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

I changed my mind it C, SecurityAdmin that inherit the priveleges of UserAdmin

upvoted 1 times

**bonniezl** 3 years, 5 months ago

The correct answer is C. The user administrator (USERADMIN) role includes the privileges to create and manage users and roles (assuming ownership of those roles or users has not been transferred to another role). The security administrator (i.e users with the SECURITYADMIN system role) role includes the global MANAGE GRANTS privilege to grant or revoke privileges on objects in the account. The USERADMIN role is a child of this role in the default access control hierarchy. https://docs.snowflake.com/en/user-guide/security-access-control-considerations

upvoted 2 times

**Omkarshingare** 3 years, 7 months ago

The user administrator (USERADMIN) role includes the privileges to create and manage users and roles (assuming ownership of those roles or users has not been transferred to another role).

upvoted 1 times

**psilvabr** 3 years, 8 months ago

B ORGADMIN to manage accounts.

upvoted 2 times

**Lavpak** 3 years, 11 months ago

An account can be created by an ORGADMIN through the web interface or using SQL https://docs.snowflake.com/en/user-guide/organizations-manage-accounts.html

upvoted 3 times

Question #185 *Topic 1*

Assume there is a table consisting of five micro-partitions with values ranging from A to Z.

Which diagram indicates a well-clustered table?




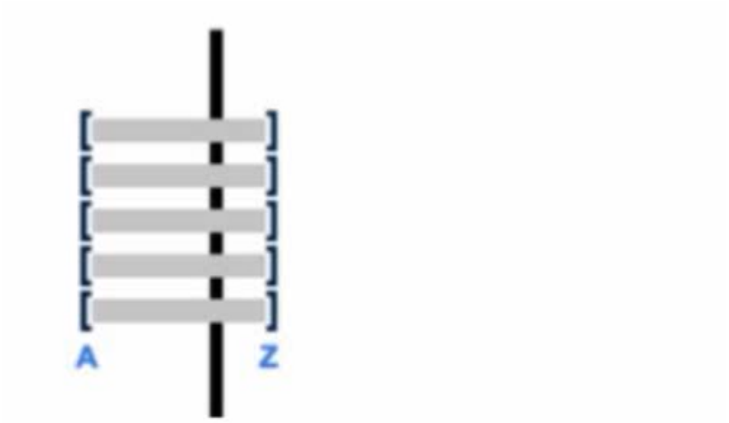


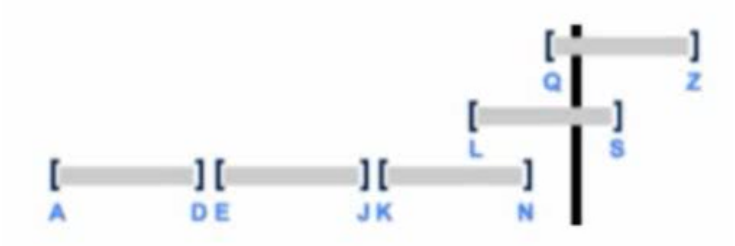





A.

B.

C.

D.

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A Is correct,It dont have overlapping

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A is correct

upvoted 1 times

**Marge23** 3 years, 1 month ago

no overlapping is a well clustered.

upvoted 1 times

**lukalooko** 3 years, 3 months ago

It is A when only going off these 4 images. However in real scenarios, these things should be kept in mind: * The clustering depth for a table is not an absolute or precise measure of whether the table is well-clustered. Ultimately, query performance is the best indicator of how well-clustered a table is: ->If queries on a table are performing as needed or expected, the table is likely well-clustered. -> If query performance degrades over time, the table is likely no longer well-clustered and may benefit from clustering. https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions The images are exactly from that link, which also states "The diagram is not intended to represent an actual table. In an actual table, with data contained in a large numbers of micro-partitions, reaching a constant state across all micro-partitions is neither likely nor required to improve query performance."

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**AbhaySAP** 3 years, 4 months ago

Was on exam 2nd July 2023

upvoted 2 times

**SV1122** 3 years, 10 months ago

Was on exam Dec 31st 2022

upvoted 2 times

**halol** 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions.html

upvoted 3 times

Question #186 *Topic 1*

What feature can be used to reorganize a very large table on one or more columns?

A. Micro-partitions

B. Clustering keys

C. Key partitions

D. Clustered partitions

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Clustering keys

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Clustering keys

upvoted 1 times

**lukalooko** 3 years, 3 months ago

Correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**halol** 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html

upvoted 3 times

Question #187 *Topic 1*

What is an advantage of using an explain plan instead of the query profiler to evaluate the performance of a query?

A. The explain plan output is available graphically.

B. An explain plan can be used to conduct performance analysis without executing a query.

C. An explain plan will handle queries with temporary tables and the query profiler will not.

D. An explain plan's output will display automatic data skew optimization information.

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**SV1122** Highly Voted 3 years, 4 months ago

EXPLAIN compiles the SQL statement, but does not execute it, so EXPLAIN does not require a running warehouse. Although EXPLAIN does not consume any compute credits, the compilation of the query does consume Cloud Service credits, just as other metadata operations do.
https://docs.snowflake.com/en/sql-reference/sql/explain.html

upvoted 11 times

**_yyukta** Most Recent 2 years, 2 months ago

B. An explain plan can be used to conduct performance analysis without executing a query.

upvoted 1 times

**lukalooko** 2 years, 9 months ago

Correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**halol** 3 years, 5 months ago

https://docs.snowflake.com/en/sql-reference/sql/explain.html

upvoted 2 times

Question #188 *Topic 1*

Which data types are supported by Snowflake when using semi-structured data? (Choose two.)

A. VARIANT

B. VARRAY

C. STRUCT

D. ARRAY

E. QUEUE

**Suggested Answer:** *A, D*

*Community vote distribution*

AD (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Array,Variant

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: AD

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A and D are correct

upvoted 1 times

**Marge23** 3 years, 1 month ago

variant , array

upvoted 1 times

**lukalooko** 3 years, 3 months ago

From documentation: https://docs.snowflake.com/en/sql-reference/data-types-semistructured.html Variant, Object, Array - so of the available options in question. A and D.

upvoted 2 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**alanhuangjob** 3 years, 7 months ago

Variant, Array and Object types.

upvoted 2 times

**SV1122** 3 years, 10 months ago

Was on exam Dec 31st 2022

upvoted 2 times

**SV1122** 3 years, 10 months ago

Variant Object Array https://docs.snowflake.com/en/sql-reference/data-types-semistructured.html

upvoted 1 times

**halol** 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/data-types-semistructured.html

upvoted 2 times

Question #189 *Topic 1*

Why does Snowflake recommend file sizes of 100-250 MB compressed when loading data?

A. Optimizes the virtual warehouse size and multi-cluster setting to economy mode

B. Allows a user to import the files in a sequential order

C. Increases the latency staging and accuracy when loading the data

D. Allows optimization of parallel operations

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**_yyukta** Most Recent 2 years, 2 months ago

D. Allows optimization of parallel operations

upvoted 1 times

**Marge23** 2 years, 7 months ago

D correct

upvoted 2 times

**lukalooko** 2 years, 9 months ago

Correct. Another thing to keep in mind in actual operation in utilizing, for example a Snowpipe, it is less costly to have lots of data in large compiled files flowing through a Snowpipe. For every 1,000 files queued, you have to pay for 0.06 credits. I have tested this myself.

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**c0d3g** 3 years, 1 month ago

Allows optimization of parallel operations https://docs.snowflake.com/en/user-guide/data-load-considerations-prepare

upvoted 1 times

**SV1122** 3 years, 4 months ago

To optimize the number of parallel operations for a load, we recommend aiming to produce data files roughly 100-250 MB (or larger) in size compressed.

upvoted 2 times

**halol** 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/data-load-considerations-prepare.html

upvoted 2 times

Question #190 *Topic 1*

Which of the following features are available with the Snowflake Enterprise edition? (Choose two.)

A. Database replication and failover

B. Automated index management

C. Customer managed keys (Tri-secret secure)

D. Extended time travel

E. Native support for geospatial data

**Suggested Answer:** *D, E*

*Community vote distribution*

DE (91%) (4 (4

**Raju039** Highly Voted 3 years, 7 months ago

A. Database replication and failover - Business B. Automated index management - No index in Snowflake C. Customer managed keys (Tri-secret secure) - Business D. Extended time travel - Enterprise E. Native support for geospatial data - Standard

upvoted 14 times

**tejarun** Highly Voted 3 years, 11 months ago

A. Database replication and fail over and C. Customer managed keys (Tri-secret secure) are supported Business Critical and above only. https://docs.snowflake.com/en/user-guide/intro-editions.html

upvoted 13 times

**MultiCloudIronMan** Most Recent 1 year, 9 months ago

ADE is correct

upvoted 1 times

**Wydad_19** 1 year, 11 months ago

Just D.

upvoted 2 times

**nehanovigo** 2 years, 3 months ago

D&E - Database replication /failover is Business critical and higher

upvoted 1 times

**Apoc2662** 2 years, 3 months ago

Actually just D. For A, data replication is included with standard edition but failover is part of business critical. https://docs.snowflake.com/en/user-guide/intro-editions

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

A and D are correct

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct answer:DE

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A and D are correct

upvoted 1 times

**sandy91** 2 years, 9 months ago

D and E are correct. Native support for geospatial data is for all the versions. Extended time travel is for Enterprise and above.

upvoted 2 times

**uadabala** 2 years, 9 months ago

CDhttps://docs.snowflake.com/en/user-guide/security-encryption-manage?utm_source=legacy&utm_medium=serp&utm_term=Customer+managed+keys

upvoted 1 times

**lol632** 2 years, 11 months ago

This is correct

upvoted 2 times

**Colourseun** 2 years, 11 months ago

The following features are available with the Snowflake Enterprise edition: C: Customer managed keys (Tri-secret secure) D: Extended time travel

upvoted 2 times

**saikot** 3 years, 2 months ago

The two features that are available with the Snowflake Enterprise edition are: DExtended time travel up to 90 days1. CCustomer managed keys (Tri-secret secure)1.

upvoted 2 times

**Amitsnowflake** 3 years, 2 months ago

Enterprise and lower

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**ajay_1233456** 3 years, 4 months ago

it will be CD

upvoted 1 times

Question #191 *Topic 1*

What is the default file size when unloading data from Snowflake using the COPY command?

A. 5 MB

B. 8 GB

C. 16 MB

D. 32 MB

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

C.16 MB

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C. 16 MB

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**dmitriypo** 3 years, 5 months ago

16 Mb https://docs.snowflake.com/en/sql-reference/sql/copy-into-location#copy-options-copyoptions

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

16 MB maximum file size

upvoted 1 times

**SV1122** 3 years, 10 months ago

By default, COPY INTO location statements separate table data into a set of output files to take advantage of parallel operations. The maximum size for each file is set using the MAX_FILE_SIZE copy option. The default value is 16777216 (16 MB) but can be increased to accommodate larger files.
https://docs.snowflake.com/en/user-guide/data-unload-considerations.html

upvoted 3 times

**halol** 3 years, 11 months ago

correct C

upvoted 1 times

Question #192 *Topic 1*

What features that are part of the Continuous Data Protection (CDP) feature set in Snowflake do not require additional configuration? (Choose two.)

A. Row level access policies

B. Data masking policies

C. Data encryption

D. Time Travel

E. External tokenization

**Suggested Answer:** *C, D*

*Community vote distribution*

CD (100%)

**halol** Highly Voted 3 years, 10 months ago

correct https://docs.snowflake.com/en/user-guide/data-cdp.html

upvoted 5 times

**kolossal** Most Recent 1 year, 12 months ago

External Tokenization is available starting from Enterprise Edition

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

Data Encryption

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C and D are correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**c0d3g** 3 years, 7 months ago

Data encryption Time Travel

upvoted 2 times

Question #193 *Topic 1*

Which Snowflake layer is always leveraged when accessing a query from the result cache?

A. Metadata

B. Data Storage

C. Compute

D. Cloud Services

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Cloud Services layer

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. Cloud Services

upvoted 1 times

**Marge23** 3 years, 1 month ago

d correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**c0d3g** 3 years, 7 months ago

Cloud Services

upvoted 1 times

**halol** 3 years, 10 months ago

correct https://community.snowflake.com/s/article/Caching-in-Snowflake-Data-Warehouse#:~:text=Snowflake%20Cache%20Layers&text=Result%20Cache%3A%20Which%20holds%20the,underlying%20data%20has%20not%20changed.

upvoted 4 times

Question #194 *Topic 1*

Which connectors are available in the downloads section of the Snowflake web interface (UI)? (Choose two.)

A. SnowSQL

B. JDBC

C. ODBC

D. HIVE

E. Scala

**Suggested Answer:** *A, C*

*Community vote distribution*




**dmitriypo** Highly Voted 3 years, 5 months ago

The question relates to the old UA. JDBC and ODBC drivers can be seen there. There are no Download section in the ne Snowsite UA.

upvoted 8 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Snowsql and odbc

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A and C are correct

upvoted 2 times

**VenkataPhaniPavanKumar** 2 years, 9 months ago

Correct Answer: AC (SNOWSQL and ODBC) https://docs.snowflake.com/en/user-guide/snowflake-client-repository

upvoted 2 times

**azlovemovie** 2 years, 10 months ago

This is based on old snowflake (the classic user interface)

upvoted 1 times

**Pradeep_7869** 3 years ago

A and C

upvoted 1 times

**Learner21525** 3 years ago

The question relates to the old UA. JDBC and ODBC drivers can be seen there. There are no Download section in the ne Snowsite UA.

upvoted 1 times

**PoornaChandraIn** 3 years, 3 months ago

AC is correct answer.. https://docs.snowflake.com/en/user-guide/snowflake-client-repository

upvoted 2 times

**MultiCloudIronMan** 3 years, 3 months ago

https://www.bing.com/search?q=Which+connectors+are+available+in+the+downloads+section+of+the+Snowflake+web+interface+%28UI%29%3F+%28Choose+two.%29&qs=n&form=QBRE&sp=-1&lq=1&pq=which+connectors+are+available+in+the+downloads+section+of+the+snowflake+web+interface+%28ui%29%3F+%28choose+two.%29&sc=1-106&sk=&cvid=4A74A11C8A8A4B7E943BE38A23305434&ghsh=0&ghacc=0&ghpl=

upvoted 2 times

**JJ_Snowfake** 3 years, 4 months ago

https://developers.snowflake.com/drivers-and-libraries/

upvoted 1 times

**valiayaan** 3 years, 4 months ago

AC is the correct answer

upvoted 1 times

**Abhishek1180** 3 years, 4 months ago

https://www.examtopics.com/discussions/snowflake/view/69294-exam-snowpro-core-topic-1-question-56-discussion/ - what kind contradiction is this?

upvoted 2 times

**bonniezl** 3 years, 5 months ago

https://developers.snowflake.com/drivers-and-libraries/

upvoted 2 times

**RajuNaik** 3 years, 6 months ago

ABC are correct

upvoted 1 times

**Dan1420** 3 years, 6 months ago

JDBC Driver Download from Maven Docs ODBC Driver Download from Snowflake SnowSQL Download from Snowflake

https://developers.snowflake.com/drivers-and-libraries/

upvoted 2 times

**Omkarshingare** 3 years, 7 months ago

SnowSQL and ODBC is correct answer

upvoted 1 times

**Raju039** 3 years, 7 months ago

https://developers.snowflake.com/drivers-and-libraries/

upvoted 2 times

Question #195 *Topic 1*

A Snowflake Administrator needs to ensure that sensitive corporate data in Snowflake tables is not visible to end users, but is partially visible to functional managers.

How can this requirement be met?

A. Use data encryption.

B. Use dynamic data masking.

C. Use secure materialized views.

D. Revoke all roles for functional managers and end users.

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Dynamic Data Masking

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Use dynamic data masking.

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Even though a secured materialised view can do the job at higher cost.

upvoted 1 times

**SV1122** 3 years, 10 months ago

Masking policy administrators can implement a masking policy such that analysts (i.e. users with the custom ANALYST role) can only view the last four digits of a phone number and none of the social security number, while customer support representatives (i.e. users with the custom SUPPORT role) can view the entire phone number and social security number for customer verification use cases. https://docs.snowflake.com/en/user-guide/security-column-intro.html#what-are-masking-policies

upvoted 3 times

**AravindhTN** 3 years, 10 months ago

Correct

upvoted 2 times

Question #196 *Topic 1*

Users are responsible for data storage costs until what occurs?

A. Data expires from Time Travel

B. Data expires from Fail-safe

C. Data is deleted from a table

D. Data is truncated from a table

**Suggested Answer:** *B*

*Community vote distribution*

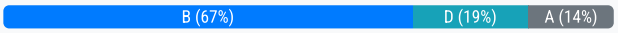


**SV1122** Highly Voted 3 years, 10 months ago

Selected Answer: B

Storage fees are incurred for maintaining historical data during both the Time Travel and Fail-safe periods.

https://docs.snowflake.com/en/user-guide/data-cdp-storage-costs.html

upvoted 7 times

**SV1122** Highly Voted 3 years, 10 months ago

Storage fees are incurred for maintaining historical data during both the Time Travel and Fail-safe periods. https://docs.snowflake.com/en/user-guide/data-cdp-storage-costs.htmlThat means answer A should be correct!?Order : data + time travel (0-90) + fail safe(7)

upvoted 7 times

**Fiko** 3 years, 9 months ago

That means answer A should be correct!?

upvoted 2 times

**Fiko** 3 years, 9 months ago

That means answer A should be correct!?Order : data + time travel (0-90) + fail safe(7)

upvoted 2 times

**Abhishek1180** 3 years, 4 months ago

Order : data + time travel (0-90) + fail safe(7)

upvoted 2 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B.Data expires from Fail-safe

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Data expires from Fail-safe

upvoted 1 times

**Mike722** 2 years, 8 months ago

B: https://docs.snowflake.com/en/user-guide/tables-storage-considerations Storage is calculated and charged for data regardless of whether it is in the Active, Time Travel, or Fail-safe state. Because these life-cycle states are sequential, updated/deleted data protected by CDP will continue to incur storage costs until the data leaves the Fail-safe state.

upvoted 2 times

**Colourseun** 2 years, 11 months ago

The answer is D. Data is truncated from a table. Once data is truncated from a table, it is permanently deleted and cannot be recovered. Snowflake charges for data storage until the data is truncated.

upvoted 4 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**srj_mehta** 3 years, 4 months ago

B IS RIGHT Storage is calculated and charged for data regardless of whether it is in the Active, Time Travel, or Fail-safe state. Because these life-cycle states are sequential, updated/deleted data protected by CDP will continue to incur storage costs until the data leaves the Fail-safe state. https://docs.snowflake.com/en/user-guide/tables-storage-considerations

upvoted 2 times

**Omkarshingare** 3 years, 7 months ago

Data files staged in Snowflake internal stages are not subject to the additional costs associated with Time Travel and Fail-safe, but they do incur standard data storage costs.

upvoted 1 times

**shyemko** 3 years, 7 months ago

No data = no costNo. If the table is truncated, time -travel comes into the picture which makes you responsible for storage costs again.

upvoted 1 times

**shyemko** 3 years, 7 months ago

Selected Answer: D

No data = no cost

upvoted 1 times

**Naren9094** 3 years, 3 months ago

No. If the table is truncated, time -travel comes into the picture which makes you responsible for storage costs again.

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

I feel D

upvoted 3 times

**KarBiswa** 3 years, 8 months ago

Actually I feel the meaning of the question is what is event that incurs extra cost, in that case the table where data modifications are more incur more cost in that case truncating is churning the entire table data will incur maximum cost so I feel the right one is D

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

Bringing data from fail safe is not free of cost

upvoted 1 times

**NabilR** 3 years, 9 months ago

Should be A, since data can remain in Time Travel for 90 days, while Fail Safe is for 7 days only.The Fail-Safe append to the 90 days, so its up to 97 days.

upvoted 2 times

**NabilR** 3 years, 9 months ago

Selected Answer: A

Should be A, since data can remain in Time Travel for 90 days, while Fail Safe is for 7 days only.

upvoted 2 times

**mlaus** 3 years, 9 months ago

The Fail-Safe append to the 90 days, so its up to 97 days.

upvoted 1 times

**halol** 3 years, 11 months ago

Storage is calculated and charged for data regardless of whether it is in the Active, Time Travel, or Fail-safe state. Because these life-cycle states are sequential, updated/deleted data protected by CDP will continue to incur storage costs until the data leaves the Fail-safe state. https://docs.snowflake.com/en/user-guide/tables-storage-considerations.html

upvoted 3 times

Question #197 *Topic 1*

A user has an application that writes a new file to a cloud storage location every 5 minutes.

What would be the MOST efficient way to get the files into Snowflake?

A. Create a task that runs a COPY INTO operation from an external stage every 5 minutes.

B. Create a task that PUTS the files in an internal stage and automate the data loading wizard.

C. Create a task that runs a GET operation to intermittently check for new files.

D. Set up cloud provider notifications on the file location and use Snowpipe with auto-ingest.

**Suggested Answer:** *D*

*Community vote distribution*

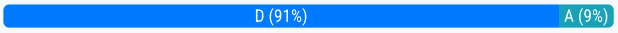


**JUHYEON** Highly Voted 3 years, 10 months ago

i think D

upvoted 5 times

**niho1993** Most Recent 1 year, 1 month ago

“With Snowpipe’s auto-ingest feature, Snowflake can be configured to listen to event notifications from cloud storage. When new data files are added, Snowpipe automatically loads them into a target table in Snowflake.”

upvoted 1 times

**37ceea2** 1 year, 8 months ago

D is not a good answer. When we have cloud provider notification, the AUTO-INGEST for SNOWPIPE has to be FALSE. A is good, due to the fact that it refers to External Stage and, in the "SNOWFLAKE language" when it is about Cloud Storage it is about files on External Stages. Hence, we have to load from External Stages to Table, using COPY INTO. Last idea, is what does the question understand by "with AUTO-INGEST"? AUTO-INGEST=TRUE or AUTO-INGEST=FALSE. I still remain to A.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

D. Set up cloud provider notifications on the file location and use Snowpipe with auto-ingest.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. Set up cloud provider notifications on the file location and use Snowpipe with auto-ingest.

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**harshkoshta** 3 years, 4 months ago

A can also right, because we can schedule task

upvoted 1 times

**c0d3g** 3 years, 7 months ago

D. Set up cloud provider notifications on the file location and use Snowpipe with auto-ingest.

upvoted 2 times

Question #198 *Topic 1*

What affects whether the query results cache can be used?

A. If the query contains a deterministic function

B. If the virtual warehouse has been suspended

C. If the referenced data in the table has changed

D. If multiple users are using the same virtual warehouse

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**SnowProCertDec22** Highly Voted 3 years, 10 months ago

answer correct is C https://community.snowflake.com/s/article/Understanding-Result-Caching

upvoted 8 times

**pranalig** Most Recent 2 years, 4 months ago

answer correct is C

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

In Snowflake, the query results cache can be used based on certain conditions. Specifically, whether the referenced data in the table has changed affects the use of the query results cache. If the data has changed since the last query, the cache cannot be used. Therefore, the correct answer is: C. If the referenced data in the table has changed Changes to the underlying data will invalidate the cached results, necessitating a new query execution to retrieve the updated data.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C. If the referenced data in the table has changed

upvoted 1 times

**Aaron_w_mendelson** 2 years, 9 months ago

Results set cache lives in the cloud services layer not the warehouse

upvoted 1 times

**BobFar** 3 years ago

Answer is C. A is wrong because non-deterministic function has impact on result cache except current_date Answer B is wrong because : If a virtual warehouse is suspended, it won't be available to execute any queries, so the result cache won't come into play for queries that would have been processed by that warehouse. When a virtual warehouse is resumed after being suspended, any query that benefits from the result cache will continue to do so. The result cache is independent of the state of the virtual warehouse. The result cache is a feature that helps optimize query performance regardless of the state (active or suspended) of the virtual warehouse. Answer D is wrong because same query returns from the cache.

upvoted 1 times

**shanksund** 3 years ago

This is 100% B, the Cache is cleared whenever a VW is suspended

upvoted 1 times

**shanksund** 3 years ago

This is 100% B, the Cache is cleared whenever a VW is suspendedOnly the local disk cache is lost when the VW is suspended. Result Cache is in the service layer and will survive (at least 24 hours) even if all warehouses are suspended.

upvoted 1 times

**Mike722** 2 years, 8 months ago

Only the local disk cache is lost when the VW is suspended. Result Cache is in the service layer and will survive (at least 24 hours) even if all warehouses are suspended.

upvoted 4 times

**Thaiya** 3 years, 1 month ago

Answer A - if count() sum() or predictable query is used, not C - if data in table is changed the saved answer in cache will become incorrectMy bad, i didnt read question properly. affects then correct answer is C.

upvoted 1 times

**Thaiya** 3 years, 1 month ago

Answer A - if count() sum() or predictable query is used, not C - if data in table is changed the saved answer in cache will become incorrect

upvoted 1 times

**Thaiya** 3 years, 1 month ago

My bad, i didnt read question properly. affects then correct answer is C.

upvoted 2 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**bonniezl** 3 years, 5 months ago

C is correct

upvoted 1 times

**mighty_kafta** 3 years, 8 months ago

C is the correct answer

upvoted 2 times

**n21007** 3 years, 8 months ago

C is the correct answer

upvoted 1 times

**muks86** 3 years, 8 months ago

Moderator should update database with correct answers,as in this case it should be C

upvoted 4 times

**KarBiswa** 3 years, 8 months ago

C is correct agree

upvoted 1 times

**JRayan** 3 years, 8 months ago

agree C is the right answer

upvoted 1 times

**jtm22** 3 years, 9 months ago

C appears correct

upvoted 1 times

**Naah** 3 years, 10 months ago

Answer: C Result Cache: Which holds the results of every query executed in the past 24 hours. These are available across virtual warehouses, so query results returned to one user is available to any other user on the system who executes the same query, provided the underlying data has not changed.

upvoted 4 times

Question #199 *Topic 1*

Which of the following is an example of an operation that can be completed without requiring compute, assuming no queries have been executed previously?

A. SELECT SUM (ORDER_AMT) FROM SALES;

B. SELECT AVG(ORDER_QTY) FROM SALES;

C. SELECT MIN(ORDER_AMT) FROM SALES;

D. SELECT ORDER_AMT * ORDER_QTY FROM SALES;

**Suggested Answer:** *C*

*Community vote distribution*

C (96%) (4

**tejarun** Highly Voted 3 years, 11 months ago

Average and Sum functions will need compute while min function can execute from cloud services layer

upvoted 14 times

**shyemko** Highly Voted 4 months ago

Selected Answer: C

Bulk data load Stored in the metadata of the target table for 64 days. Available upon completion of the COPY statement as the statement output. Snowpipe Stored in the metadata of the pipe for 14 days. Must be requested from Snowflake via a REST endpoint, SQL table function, or ACCOUNT_USAGE view. source : https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro

upvoted 7 times

**ntechkuma** 3 years, 8 months ago

Moderator should update correct answers. most of the questions have wrong answers.+YES AGREE

upvoted 7 times

**EmiB** 4 months ago

Selected Answer: C

COPY INTO - Stored in the metadata of the target table for 64 days SNOWPIPE - Stored in the metadata of the pipe for 14 days

upvoted 4 times

**VLMBL_94** 2 years, 2 months ago

+YES AGREE

upvoted 1 times

**_yyukta** 4 months ago

Selected Answer: C

C. 14 days

upvoted 2 times

**jhfather** 1 year, 8 months ago

D: SELECT ORDER_AMT * ORDER_QTY FROM SALES involves a simple operation of multiplying the values of two columns. This operation can be performed solely by scanning the data and does not require a Compute Engine.

upvoted 1 times

**Naah** 4 months ago

C https://docs.snowflake.com/en/sql-reference/functions/copy_history.html

upvoted 2 times

**vivekmani2021** 2 years, 3 months ago

c , Min

upvoted 1 times

**Rob__C** Most Recent 4 months ago

https://docs.snowflake.com/en/sql-reference/info-schema/load_history.html

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

C is correct answer

upvoted 1 times

**YoshiSamoZou** 2 years, 4 months ago

c correct

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct answer:C

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

answer is C ,min and max can execute from cloud service

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C. SELECT MIN(ORDER_AMT) FROM SALES;

upvoted 1 times

**Gaurangvp44** 3 years ago

MIN Is t h e right

upvoted 1 times

**Marge23** 3 years, 1 month ago

MIN no need vw

upvoted 1 times

**singhks** 3 years, 2 months ago

C is the answer

upvoted 1 times

**bonniezl** 3 years, 5 months ago

C is correct

upvoted 1 times

**shyemko** 3 years, 7 months ago

Below are what is collected in Metadata : - Range of values including Max, Min & Count - Number of distinct values -Other optimization metadata

upvoted 1 times

**Josa2** 3 years, 8 months ago

I believe it's C, since once the METADATA stores it, it will know what's the DETERMINISTIC value without having to perform any calculation as suggested by A and B or run over all micro-partitions as it would occur in D.

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

MIN , MAX is ideal for this case

upvoted 1 times

**thejeby** 3 years, 9 months ago

C is the answer

upvoted 1 times

Question #200 *Topic 1*

How many days is load history for Snowpipe retained?

A. 1 day

B. 7 days

C. 14 days

D. 64 days

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**shyemko** Highly Voted 3 years, 7 months ago

Bulk data load Stored in the metadata of the target table for 64 days. Available upon completion of the COPY statement as the statement output. Snowpipe Stored in the metadata of the pipe for 14 days. Must be requested from Snowflake via a REST endpoint, SQL table function, or ACCOUNT_USAGE view. source : https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro

upvoted 7 times

**MultiCloudIronMan** Most Recent 1 year, 9 months ago

Agreed

upvoted 1 times

**dongwenlong3** 1 year, 11 months ago

C.14 days

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

C.14 days

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C. 14 days

upvoted 2 times

**Colourseun** 2 years, 11 months ago

Load history for Snowpipe is retained for 14 days:C Snowflake stores load history for Snowpipe in the metadata of the pipe object. This history includes information such as the date and time of the load, the number of files loaded, and the size of the data loaded. SELECT * FROM INFORMATION_SCHEMA.PIPE_LOAD_HISTORY WHERE PIPE_NAME = ' ' ORDER BY START_TIME DESC;

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**EmiB** 3 years, 8 months ago

COPY INTO - Stored in the metadata of the target table for 64 days SNOWPIPE - Stored in the metadata of the pipe for 14 days

upvoted 4 times

**muks86** 3 years, 8 months ago

Correct ans 14 days

upvoted 1 times

**Naah** 3 years, 10 months ago

C https://docs.snowflake.com/en/sql-reference/functions/copy_history.html

upvoted 2 times

**Rob__C** 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/info-schema/load_history.html

upvoted 2 times

Question #201 *Topic 1*

What Snowflake features allow virtual warehouses to handle high concurrency workloads? (Choose two.)

A. The ability to scale up warehouses

B. The use of warehouse auto scaling

C. The ability to resize warehouses

D. Use of multi-clustered warehouses

E. The use of warehouse indexing

**Suggested Answer:** *B, D*

*Community vote distribution*

BD (100%)

**Jiaa** Highly Voted 3 years, 4 months ago

I think BD https://docs.snowflake.com/en/sql-reference/parameters.html#max-concurrency-level https://docs.snowflake.com/en/sql-reference/parameters.html#max-concurrency-level

upvoted 5 times

**_yyukta** Most Recent 2 years, 2 months ago

B and D are correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**dmitriypo** 2 years, 12 months ago

Definitely BD

upvoted 2 times

**shyemko** 3 years, 1 month ago

Auto-scaling Multicluster warehouse

upvoted 2 times

Question #202 *Topic 1*

Which COPY INTO command outputs the data into one file?

A. SINGLE=TRUE

B. MAX_FILE_NUMBER=1

C. FILE_NUMBER=1

D. MULTIPLE=FALSE

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**halol** Highly Voted 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/sql/copy-into-location.html

upvoted 5 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A is correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. SINGLE=TRUE are correct

upvoted 1 times

**shyemko** 3 years, 7 months ago

SINGLE = TRUE

upvoted 1 times

**EmiB** 3 years, 7 months ago

A. SINGLE=TRUE

upvoted 1 times

**vmsaran** 3 years, 8 months ago

Answer is A

upvoted 1 times

Question #203 *Topic 1*

In which scenarios would a user have to pay Cloud Services costs? (Choose two.)

A. Compute Credits = 50 Credits Cloud Services = 10

B. Compute Credits = 80 Credits Cloud Services = 5

C. Compute Credits = 100 Credits Cloud Services = 9

D. Compute Credits = 120 Credits Cloud Services = 10

E. Compute Credits = 200 Credits Cloud Services = 26

**Suggested Answer:** *A, E*

*Community vote distribution*

AE (100%)

**tejarun** Highly Voted 3 years, 11 months ago

(Credits Cloud Services / Compute Credits)% > 10% for A and E

upvoted 21 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A and E are Correct

upvoted 1 times

**gunasekarantjk1996** 2 years, 5 months ago

AE In Snowflake, a user would have to pay for Cloud Services costs in the following scenarios: Compute Credits = 50 Credits Cloud Services = 10 (A): Usage of Cloud Services is charged only if the daily consumption of cloud services exceeds 10% of the daily usage of virtual warehouses1. In this case, Cloud Services usage is 20% of Compute Credits, which exceeds the 10% threshold1. Compute Credits = 200 Credits Cloud Services = 26 (E): In this case, Cloud Services usage is 13% of Compute Credits, which also exceeds the 10% threshold1. Please note that if Cloud Services consumption is less than or equal to 10% of warehouse compute credits on a given day, no additional charges are incurred1. Therefore, scenarios B, C, and D would not incur additional Cloud Services costs1.

upvoted 4 times

**_yyukta** 2 years, 8 months ago

AE are correct

upvoted 1 times

**0e504b5** 2 years, 9 months ago

AE. Examples are here: https://docs.snowflake.com/en/user-guide/cost-understanding-compute

upvoted 2 times

**Marge23** 3 years, 1 month ago

over 10% usage of the credit

upvoted 2 times

**bbrandhh** 3 years, 3 months ago

What you can also do for this problem is you can multiply the credits cloud service by itself and see if that number is greater than the credits So for example for A. 90... 10 * 10 = 100 > 90 so A would be correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Thats the answer please can Nitin001 our Maths major confirm

upvoted 1 times

**bonniezl** 3 years, 5 months ago

AE is correct

upvoted 1 times

**Nitin001** 3 years, 5 months ago

CE right? Cloud service cost is charged to User only when Cloud service cost exceeds compute cost by more than 10%. A: 10/50 = 10% (did not exceed 10%) B: 5/80 = 6.25% (less than 10%) C: 9/10 = 90% (more than 10%) D: 10/120 = 8.33% (less than 10%) E: 26/200 = 13% (more than 10%) I am a math major, really wondering what is up with math skills of everyone!? Or am I missing something here!Correction: ACE right? Cloud service cost is charged to User only when Cloud service cost exceeds compute cost by more than 10%. A: 10/50 = 20% (more than 10%) B: 5/80 = 6.25% (less than 10%) C: 9/10 = 90% (more than 10%) D: 10/120 = 8.33% (less than 10%) E: 26/200 = 13% (more than 10%)9/100, not 9/10For option C, you took the wrong values, it is 9/100 = 9%

upvoted 1 times

**Nitin001** 3 years, 5 months ago

CE right? Cloud service cost is charged to User only when Cloud service cost exceeds compute cost by more than 10%. A: 10/50 = 10% (did not exceed 10%) B: 5/80 = 6.25% (less than 10%) C: 9/10 = 90% (more than 10%) D: 10/120 = 8.33% (less than 10%) E: 26/200 = 13% (more than 10%) I am a math major, really wondering what is up with math skills of everyone!? Or am I missing something here!

upvoted 1 times

**Nitin001** 3 years, 5 months ago

Correction: ACE right? Cloud service cost is charged to User only when Cloud service cost exceeds compute cost by more than 10%. A: 10/50 = 20% (more than 10%) B: 5/80 = 6.25% (less than 10%) C: 9/10 = 90% (more than 10%) D: 10/120 = 8.33% (less than 10%) E: 26/200 = 13% (more than 10%)

upvoted 2 times

**Nitin001** 3 years, 5 months ago

Correction: ACE right? Cloud service cost is charged to User only when Cloud service cost exceeds compute cost by more than 10%. A: 10/50 = 20% (more than 10%) B: 5/80 = 6.25% (less than 10%) C: 9/10 = 90% (more than 10%) D: 10/120 = 8.33% (less than 10%) E: 26/200 = 13% (more than 10%)9/100, not 9/10

upvoted 2 times

**bar_ser** 3 years, 5 months ago

9/100, not 9/10

upvoted 4 times

**sri_ch** 3 years, 5 months ago

For option C, you took the wrong values, it is 9/100 = 9%

upvoted 2 times

**shyemko** 3 years, 7 months ago

Usage for cloud services is charged only if the daily consumption of cloud services exceeds 10% of the daily usage of virtual warehouses. The charge is calculated daily (in the UTC time zone). This ensures that the 10% adjustment is accurately applied each day, at the credit price for that day.
https://docs.snowflake.com/en/user-guide/cost-understanding-compute#cloud-service-credit-usage

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

exceeded 10%

upvoted 1 times

**mlaus** 3 years, 9 months ago

(Credits Cloud Services / Compute Credits)% > 10% for A and E

upvoted 2 times

**iamPhD** 3 years, 10 months ago

A and E daily consumption exceeds 10%

upvoted 2 times

**miadiesel** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/cost-understanding-compute.html#cloud-service-credit-usage **Usage for cloud services is charged only if the daily consumption of cloud services exceeds 10% of the daily usage of virtual warehouses

upvoted 2 times

**aalonso** 3 years, 10 months ago

When cloud services credits exceeds 10%

upvoted 2 times

Question #204 *Topic 1*

A user created a new worksheet within the Snowsight UI and wants to share this with teammates.

How can this worksheet be shared?

A. Create a zero-copy clone of the worksheet and grant permissions to teammates.

B. Create a private Data Exchange so that any teammate can use the worksheet.

C. Share the worksheet with teammates within Snowsight.

D. Create a database and grant all permissions to teammates.

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**tejarun** Highly Voted 3 years, 5 months ago
https://docs.snowflake.com/en/user-guide/ui-snowsight-worksheets-gs.html#sharing-worksheets-and-folders
upvoted 9 times

**_yyukta** Most Recent 2 years, 2 months ago
C is correct
upvoted 1 times

**vmadana** 2 years, 3 months ago
https://docs.snowflake.com/en/user-guide/ui-snowsight-worksheets-gs.html#sharing-worksheets-and-folders
upvoted 2 times

**Kuma007** 2 years, 7 months ago
Select Answer : C
upvoted 2 times

**MultiCloudIronMan** 2 years, 9 months ago
Correct
upvoted 2 times

**jollysachin** 2 years, 10 months ago
https://docs.snowflake.com/en/user-guide/ui-snowsight-worksheets-gs.html#sharing-worksheets-and-folders
upvoted 2 times

**RajuNaik** 3 years ago
C is right
upvoted 2 times

**bibisa** 3 years, 1 month ago
Within Snowsight, a user can share a worksheet with teammates by clicking on the "Share" button in the toolbar and then entering the email addresses of the teammates to share the worksheet with. This will grant the specified users view access to the worksheet. A is incorrect as zero-copy clones are used for making a copy of the data that the worksheet is working on, not for sharing the worksheet itself. B is incorrect as private Data Exchange is used for sharing data sets, not worksheets. D is incorrect as creating a database and granting all permissions is not necessary for sharing a worksheet, and granting all permissions is not recommended for security reasons.
upvoted 3 times

**shyemko** 3 years, 1 month ago
Snowsight lets you share worksheets and folders with other Snowflake users in your account so that they can view or run worksheets and folders you own. You can only share directly with users that have signed into Snowsight before https://docs.snowflake.com/en/user-guide/ui-snowsight-worksheets-gs
upvoted 1 times

**psilvabr** 3 years, 1 month ago
C you can share worksheets through Snowsight.
upvoted 1 times

**Josa2** 3 years, 2 months ago

Discussions are the greatest source in this portal at least for Snowflake cert... thanks for always coloborating... and someone should ask the moderator to curate better the answers, there're several misleadingsI totally agree. Far better than over hyped Exam Labs which has tons of wrong answers.I really do not rely on the answers of the moderators here but on the answers that stem from the Discussions. The Discussion section is indeed a saving grace.

upvoted 2 times

**Josa2** 3 years, 2 months ago

Discussions are the greatest source in this portal at least for Snowflake cert... thanks for always coloborating... and someone should ask the moderator to curate better the answers, there're several misleadings

upvoted 2 times

**SteveCrown** 2 years, 11 months ago

I totally agree. Far better than over hyped Exam Labs which has tons of wrong answers.

upvoted 1 times

**SteveCrown** 2 years, 11 months ago

I totally agree. Far better than over hyped Exam Labs which has tons of wrong answers.I really do not rely on the answers of the moderators here but on the answers that stem from the Discussions. The Discussion section is indeed a saving grace.

upvoted 1 times

**SteveCrown** 2 years, 11 months ago

I really do not rely on the answers of the moderators here but on the answers that stem from the Discussions. The Discussion section is indeed a saving grace.

upvoted 3 times

**KarBiswa** 3 years, 2 months ago

C correct

upvoted 1 times

**mlaus** 3 years, 3 months ago

C is right

upvoted 1 times

**iamPhD** 3 years, 4 months ago

C is right

upvoted 1 times

**aalonso** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/ui-snowsight-worksheets-gs.html#sharing-worksheets-and-folders

upvoted 1 times

Question #205 *Topic 1*

How can a row access policy be applied to a table or a view? (Choose two.)

A. Within the policy DDL

B. Within the create table or create view DDL

C. By future APPLY for all objects in a schema

D. Within a control table

E. Using the command ALTER [object] ADD ROW ACCESS POLICY [policy];

**Suggested Answer:** *B, E*

*Community vote distribution*




**tejarun** Highly Voted 3 years, 11 months ago

B, E https://docs.snowflake.com/en/user-guide/security-row-intro.html#apply-a-row-access-policy-to-a-table-or-view

upvoted 19 times

**_S_A_M_** Most Recent 2 years, 3 months ago

B,E IS CORRECT ANS ..ACCORDING TO SNOWFLAKE DOCUMENTATION

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B and E are correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

BE are correct

upvoted 1 times

**Avi0407** 3 years ago

Table level is the key

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**weuz** 3 years, 7 months ago

A & E are the correct choices

upvoted 1 times

**OTE** 3 years, 7 months ago

Create / Alter Table/View DDL with row access policy

upvoted 2 times

**EmiB** 3 years, 7 months ago

Selected Answer: AE

A. Within the policy DDL - CREATE TABLE t .. WITH ROW ACCESS POLICY p ON (c); E. Using the command ALTER TABLE t ADD ROW ACCESS POLICY p ON (c);

upvoted 2 times

**EmiB** 3 years, 7 months ago

A. Within the policy DDL - CREATE TABLE t .. WITH ROW ACCESS POLICY p ON (c); E. Using the command ALTER TABLE t ADD ROW ACCESS POLICY p ON (c);that is option B and E ryt

upvoted 2 times

**Omkarshingare** 3 years, 7 months ago

that is option B and E ryt

upvoted 2 times

Question #206 *Topic 1*

Which command can be used to load data files into a Snowflake stage?

A. JOIN

B. COPY INTO

C. PUT

D. GET

**Suggested Answer:** *C*

*Community vote distribution*

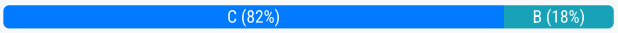


**EmiB** Highly Voted 3 years, 7 months ago

(up)load to stage - PUT (up)load from stage - COPY INTO FROM @ unload to stage - COPY INTO @ download from stage - GET

upvoted 29 times

**Geera** Highly Voted 3 years, 9 months ago

load/unload -> copy into upload/download-> put/get

upvoted 6 times

**Shipra123** Most Recent 1 year, 11 months ago

Download local files from stage = GET Upload local files to stage = PUT

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

C.PUT is correct answer

upvoted 1 times

**nexerSnow** 2 years, 5 months ago

B is the correct answer for external stage and c is the correct answer for internal stages

upvoted 2 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: C

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C. PUT

upvoted 1 times

**0e504b5** 2 years, 9 months ago

If we load from local file system, then PUT If we unload data from a Snowflake table into stage, then COPY INTO In this case, we are loading data FILES, we're not loading data from a Snowflake table into a stage. Therefore, the right answer is PUT https://docs.snowflake.com/en/user-guide/data-load-local-file-system-stage https://docs.snowflake.com/en/user-guide/data-unload-snowflake

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**chizo** 3 years, 7 months ago

C Is used to load data into a stage

upvoted 1 times

**vmsaran** 3 years, 8 months ago

C is correct

upvoted 1 times

**halol** 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/data-load-local-file-system-stage.html

upvoted 4 times

Question #207 *Topic 1*

What types of data listings are available in the Snowflake Data Marketplace? (Choose two.)

A. Reader

B. Consumer

C. Vendor

D. Standard

E. Personalized

**Suggested Answer:** *D, E*

*Community vote distribution*

DE (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Standard and Personalized Listings

upvoted 1 times

**_yyukta** 2 years, 8 months ago

DE are correct

upvoted 1 times

**Rajivnb** 3 years ago

The Snowflake Data Marketplace offers two types of data listings, Standard Data (Free) Listings and Personalized (Private)Data Listings. The key difference between standard and personalized listings is the ability to access the data share. Personalized listings allow you to control which consumers can access the data. A standard listing provides instant access to a published data set. To access the data consumers need to click Get Data button which allows you to create a Database with this dataset and will be added to Databases available. This can be done only by the Account Admin role. A personalized listing allows customers to request specific data sets. To access data from a personalized listing, consumers must submit a request by clicking the Request button, and provide their contact information. Once a request is submitted, the data provider is notified. The provider then contacts the consumer after which a personalized data set is created and shared with the consumer.STANDARD is not a name specifically called out by Snowflake. But any basic functionality with free option is considered "Standard".

upvoted 4 times

**Rajivnb** 3 years ago

Selected Answer: DE

The Snowflake Data Marketplace offers two types of data listings, Standard Data (Free) Listings and Personalized (Private)Data Listings. The key difference between standard and personalized listings is the ability to access the data share. Personalized listings allow you to control which consumers can access the data.

A standard listing provides instant access to a published data set. To access the data consumers need to click Get Data button which allows you to create a Database with this dataset and will be added to Databases available. This can be done only by the Account Admin role.

A personalized listing allows customers to request specific data sets. To access data from a personalized listing, consumers must submit a request by clicking the Request button, and provide their contact information. Once a request is submitted, the data provider is notified. The provider then contacts the consumer after which a personalized data set is created and shared with the consumer.

upvoted 3 times

**Rajivnb** 3 years ago

STANDARD is not a name specifically called out by Snowflake. But any basic functionality with free option is considered "Standard".

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

https://www.bing.com/search?pglt=41&q=What+types+of+data+listings+are+available+in+the+Snowflake+Data+Marketplace%3F&cvid=cd2c6f4d98264386b6b9dd2e423468a1&aqs=edge..69i57j69i11004.2174j0j1&FORM=ANNAB1&PC=U531

upvoted 1 times

**srj_mehta** 3 years, 4 months ago

only E https://other-docs.snowflake.com/en/collaboration/collaboration-listings-about#types-of-listings

upvoted 1 times

**muks86** 3 years, 8 months ago

The Snowflake Data Marketplace offers two types of data listings, Standard Data Listings and Personalized Data Listings

upvoted 2 times

**sonu10906997** 3 years, 9 months ago

Selected Answer: E (other options don't match) Marketplace listings can be free listings, paid listings, or personalized listings tailored to specific consumers. https://other-docs.snowflake.com/en/collaboration/collaboration-listings-about.html#types-of-listings Under MarketPlace Listing

upvoted 3 times

**BigDataBB** 3 years, 9 months ago

by this doc: https://www.linkedin.com/pulse/snowflake-data-marketplace-minzhen-yang?trk=articles_directory The right answer id D,E

upvoted 1 times

**BigDataBB** 3 years, 9 months ago

I don't find an official name as standard for the listing what I found on the documentation is: The pricing status of each data listing is indicated on the listing: Free, Paid, or Purchased. "https://other-docs.snowflake.com/en/collaboration/consumer-listings-exploring.html"

upvoted 1 times

**stornati** 3 years, 11 months ago

The Snowflake Marketplace offers two types of data listings which define how data is shared and consumed. Once published, both types of listings are displayed to consumers in the selected regions. The key difference between free and personalized listings is the ability to access the data share. Introduction to the Snowflake Marketplacehttps://other-docs.snowflake.com › marketplace › intro

upvoted 3 times

Question #208 *Topic 1*

What is the maximum Time Travel retention period for a temporary Snowflake table?

A. 90 days

B. 1 day

C. 7 days

D. 45 days

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

1 Day is correct answer

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer: B

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. 1 day

upvoted 1 times

**Marge23** 3 years, 1 month ago

1day only

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**c0d3g** 3 years, 5 months ago

Max time travel for thttps://www.examtopics.com/exams/snowflake/snowpro-core/view/21/#emporary table is one day.

upvoted 1 times

**Mubin_Sikilkar** 3 years, 9 months ago

The time travel for temporary is 0-1 days

upvoted 1 times

Question #209 *Topic 1*

When should a multi-cluster warehouse be used in auto-scaling mode?

A. When it is unknown how much compute power is needed

B. If the select statement contains a large number of temporary tables or Common Table Expressions (CTEs)

C. If the runtime of the executed query is very slow

D. When a large number of concurrent queries are run on the same warehouse

**Suggested Answer:** *D*

*Community vote distribution*

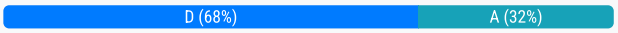


**halol** Highly Voted 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/warehouses-multicluster.html#:~:text=Auto%2Dscale,-You%20can%20choose&text=This%20mode%20is%20enabled%20by,while%20the%20warehouse%20is%20running.

upvoted 8 times

**rubiteb** 1 year, 12 months ago

Auto-Scaling can be both scaling up or down depending on the current load. Option D only describes scaling up. And if you already know what the load gonna be like then you don't need Auto-Scaling, you just set the number of cluster for your warehouse. Option A makes sense to me as this is what Auto-Scaling is for so you don't have to manually scale up or down.Correction: Auto-Scaling can be both scaling out or in. Option D only describes scaling out.

upvoted 1 times

**rubiteb** 1 year, 12 months ago

Auto-Scaling can be both scaling up or down depending on the current load. Option D only describes scaling up. And if you already know what the load gonna be like then you don't need Auto-Scaling, you just set the number of cluster for your warehouse. Option A makes sense to me as this is what Auto-Scaling is for so you don't have to manually scale up or down.

upvoted 1 times

**rubiteb** Most Recent 1 year, 12 months ago

Correction: Auto-Scaling can be both scaling out or in. Option D only describes scaling out.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

D is correct answer

upvoted 1 times

**Marty1152** 2 years, 4 months ago

In Auto-scale mode, a multi-cluster warehouse eliminates the need for resizing the warehouse or starting and stopping additional warehouses to handle fluctuating workloads. Snowflake automatically starts and stops additional clusters as needed.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. When a large number of concurrent queries are run on the same warehouse

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

correct

upvoted 1 times

**c0d3g** 3 years, 5 months ago

Auto-scale This mode is enabled by specifying different values for maximum and minimum number of clusters. In this mode, Snowflake starts and stops clusters as needed to dynamically manage the load on the warehouse: As the number of concurrent user sessions and/or queries for the warehouse increases, and queries start to queue due to insufficient resources, Snowflake automatically starts additional clusters, up to the maximum number defined for the warehouse. Similarly, as the load on the warehouse decreases, Snowflake automatically shuts down clusters to reduce the number of running clusters and, correspondingly, the number of credits used by the warehouse.

upvoted 1 times

**ntechkuma** 3 years, 8 months ago

multi cluster for concurrency

upvoted 3 times

**Geera** 3 years, 9 months ago

It can't be A because SF recommends experimenting when you don't know the right size/number of clusters for your wh.Also, compute power -> scaling up instead of scaling out

upvoted 1 times

**Geera** 3 years, 9 months ago

**Selected Answer: D**

It can't be A because SF recommends experimenting when you don't know the right size/number of clusters for your wh.

upvoted 1 times

**Geera** 3 years, 9 months ago

Also, compute power -> scaling up instead of scaling out

upvoted 2 times

**Apo1021** 3 years, 10 months ago

I think A

upvoted 2 times

**tejarun** 3 years, 11 months ago

For variable compute needs, auto scaling should be used

upvoted 4 times

Question #210 *Topic 1*

What happens when a cloned table is replicated to a secondary database? (Choose two.)

A. A read-only copy of the cloned tables is stored.

B. The replication will not be successful.

C. The physical data is replicated.

D. Additional costs for storage are charged to a secondary account.

E. Metadata pointers to cloned tables are replicated.

**Suggested Answer:** *C, D*

*Community vote distribution*

CD (88%) CE (12%)

**SV1122** Highly Voted 3 years, 10 months ago

When a cloned table is replicated to a secondary database, the physical data is also replicated, increasing the data storage usage for your account. https://docs.snowflake.com/en/user-guide/database-replication-considerations.html#:~:text=Replication%20and%20Cloning,-Cloned%20objects%20are&text=However%2C%20when%20a%20cloned%20table,storage%20usage%20for%20your%20account.

upvoted 8 times

**PrasantSadangi** Most Recent 1 year, 7 months ago

Historically Cloned objects were replicated physically rather than logically to secondary databases. That is, cloned tables in a standard database don't contribute to the overall data storage unless or until DML operations on the clone add to or modify existing data. However, when a cloned table is replicated to a secondary database, the physical data is also replicated, increasing the data storage usage for your account.

upvoted 1 times

**Jarturo** 2 years, 3 months ago

C and D https://docs.snowflake.com/en/user-guide/account-replication-considerations#replication-and-cloning

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

A and E are correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

CD are correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Vinit_Agarwal** 3 years, 4 months ago

Answers C&D are correct. For C read this: Cloned objects are replicated physically rather than logically to secondary databases. That is, cloned tables in a standard database do not contribute to the overall data storage unless or until DML operations on the clone add to or modify existing data. However, when a cloned table is replicated to a secondary database, the physical data is also replicated, increasing the data storage usage for your account. https://docs.snowflake.com/en/user-guide/account-replication-considerations#replication-and-cloni For D , read this : Charges based on replication are divided into two categories: data transfer and compute resources. Both categories are billed on the target account (i.e. the account that stores the secondary database or secondary replication/failover group that is refreshed). https://docs.snowflake.com/en/user-guide/account-replication-cost

upvoted 3 times

**c0d3g** 3 years, 5 months ago

Replication and Cloning Cloned objects are replicated physically rather than logically to secondary databases. That is, cloned tables in a standard database do not contribute to the overall data storage unless or until DML operations on the clone add to or modify existing data. However, when a cloned table is replicated to a secondary database, the physical data is also replicated, increasing the data storage usage for your account. https://docs.snowflake.com/en/user-guide/account-replication-considerations#replication-and-cloni

upvoted 1 times

**SatyamKishore** 3 years, 5 months ago

Options are not Clear . Going through the link - https://docs.snowflake.com/en/user-guide/account-replication-considerations#replication-and-cloning C is certainly correct , D if we remove secondary from the option

upvoted 2 times

**DataEngineer7331** 3 years, 6 months ago

Option A, "A read-only copy of the cloned tables is stored," is not entirely correct because the data is replicated, not just stored. Option B, "The replication will not be successful," is not correct because the replication will be successful if it is configured correctly. Option D, "Additional costs for storage are charged to a secondary account," is not correct because the cost of storage is charged to the primary account, not the secondary account.

upvoted 2 times

**harshagc180** 3 years, 11 months ago

Answer is correct

upvoted 2 times

Question #211 *Topic 1*

Snowflake supports the use of external stages with which cloud platforms? (Choose three.)

A. Amazon Web Services

B. Docker

C. IBM Cloud

D. Microsoft Azure Cloud

E. Google Cloud Platform

F. Oracle Cloud

**Suggested Answer:** *A, D, E*

*Community vote distribution*

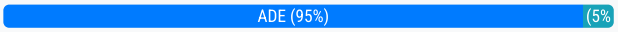


**_yyukta** 2 years, 2 months ago

ADE are correct

upvoted 1 times

**tejarun** Highly Voted 4 months ago

Selected Answer: D

Materialized view can be joined with other tables. But you cannot include JOIN in a materialized view definition

upvoted 7 times

**0e504b5** 2 years, 3 months ago

AWS, Azure and GCP

upvoted 1 times

**Josa2** Highly Voted 4 months ago

Selected Answer: D

D, how come that som many answers from basic definitions are wrong?

upvoted 5 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**lol632** 4 months ago

Selected Answer: C, D

Answer C does not make any sense at all. There is no concept of joining materialized views by another tables. Tables, per se, can join. So answer is D

upvoted 2 times

**soumavag177** 2 years, 11 months ago

These suggested answers are trying to fail us . Docker ? Seriously ? AWS, Azure , GCP (A,D,E)

upvoted 4 times

**SexyPotato** 4 months ago

Selected Answer: D

D is the answer, but B is also correct, in that you can only reference 1 table.

upvoted 2 times

**c0d3g** 2 years, 11 months ago

Amazon Web Services Microsoft Azure Cloud Google Cloud Platform

upvoted 2 times

**AS314** Most Recent 4 months ago

Selected Answer: D

https://docs.snowflake.com/en/user-guide/views-materialized.html#limitations-on-creating-materialized-views

upvoted 2 times

**Aniket_Shah** 2 years, 11 months ago

A,D,E is the right answer

upvoted 1 times

**qaiserlatif** 3 years, 1 month ago

AWS, azure and GCP

upvoted 1 times

**cgapperi** 3 years, 2 months ago

Hard to believe this one is wrong in the key.

upvoted 3 times

**muks86** 3 years, 2 months ago

The most basic question of snowflake AWS,gcp,azure

upvoted 2 times

**Varma_Saraswathula** 3 years, 2 months ago

AWS, azure and GCP

upvoted 1 times

**daddygeetha** 3 years, 4 months ago

ADE ARE THE CORRECT ANSWERS

upvoted 4 times

**aalonso** 3 years, 4 months ago

AWS, azure and GCP

upvoted 4 times

**Lavpak** 3 years, 5 months ago

ADE are correct answers

upvoted 4 times

Question #212 *Topic 1*

What is a limitation of a Materialized View?

A. A Materialized View cannot support any aggregate functions

B. A Materialized View can only reference up to two tables

C. A Materialized View cannot be joined with other tables

D. A Materialized View cannot be defined with a JOIN

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**tejarun** Highly Voted 3 years, 11 months ago

Materialized view can be joined with other tables. But you cannot include JOIN in a materialized view definition

upvoted 7 times

**Josa2** Highly Voted 3 years, 8 months ago

D, how come that som many answers from basic definitions are wrong?https://docs.snowflake.com/en/user-guide/views-materialized, can be used with a Join is a clear statement here.sorry i read the option wrong. this is correct

upvoted 5 times

**Josa2** 3 years, 8 months ago

Selected Answer: D

D, how come that som many answers from basic definitions are wrong?

upvoted 5 times

**Thaiya** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/views-materialized, can be used with a Join is a clear statement here.

upvoted 1 times

**Thaiya** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/views-materialized, can be used with a Join is a clear statement here.sorry i read the option wrong. this is correct

upvoted 1 times

**Thaiya** 3 years, 1 month ago

sorry i read the option wrong. this is correct

upvoted 1 times

**Altria** Most Recent 1 year, 1 month ago

No JOINs allowed in materialized view definition Can JOIN materialized views with other tables when querying them Single table only - cannot reference multiple tables Limited aggregate functions (COUNT, SUM, MIN, MAX) No window functions, subqueries, or complex SQL constructs Can filter, group, and perform basic calculations

upvoted 1 times

**hotflake** 2 years, 2 months ago

A. A Materialized View cannot support any aggregate functions - Wrong - supports few aggregate functions but with some restrictions. B. A Materialized View can only reference up to two tables - Wrong - can query only a single table. C. A Materialized View cannot be joined with other tables - Wrong - We can. D. A Materialized View cannot be defined with a JOIN - Correct - Joins, including self-joins, are not supported.

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

D is correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D is correct

upvoted 1 times

**DavidMiller** 2 years, 10 months ago

answer is. D

upvoted 1 times

**Thaiya** 3 years, 1 month ago

Answer is B. Query results contain a small number of rows and/or columns relative to the base table (the table on which the view is defined).The base table, only one table.

upvoted 1 times

**lol632** 2 years, 11 months ago

Answer C does not make any sense at all. There is no concept of joining materialized views by another tables. Tables, per se, can join. So answer is D

upvoted 2 times

**Thaiya** 3 years, 1 month ago

Answer is B. Query results contain a small number of rows and/or columns relative to the base table (the table on which the view is defined).The base table, only one table.D is correct, can be used with Join.

upvoted 1 times

**Thaiya** 3 years, 1 month ago

D is correct, can be used with Join.

upvoted 1 times

**jackjack123** 3 years, 3 months ago

@SexyPotato option B says 2 tables. Only D is correct.

upvoted 1 times

**SexyPotato** 3 years, 8 months ago

D is the answer, but B is also correct, in that you can only reference 1 table.

upvoted 2 times

**AS314** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/views-materialized.html#limitations-on-creating-materialized-views

upvoted 2 times

**db1290** 3 years, 11 months ago

Answer is D

upvoted 2 times

Question #213 *Topic 1*

In the Snowflake access control model, which entity owns an object by default?

A. The user who created the object

B. The SYSADMIN role

C. Ownership depends on the type of object

D. The role used to create the object

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

D is correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. The role used to create the object

upvoted 1 times

**Marge23** 3 years, 1 month ago

correct

upvoted 1 times

**Marge23** 3 years, 1 month ago

the role who created the object, owned

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**c0d3g** 3 years, 5 months ago

To own an object means that a role has the OWNERSHIP privilege on the object. Each securable object is owned by a single role, which by default is the role used to create the object.

upvoted 2 times

**EmiB** 3 years, 8 months ago

Only ROLE can OWN the object.

upvoted 2 times

**Mubin_Sikilkar** 3 years, 9 months ago

D is correct

upvoted 1 times

Question #214 *Topic 1*

What is the minimum Snowflake edition required to use Dynamic Data Masking?

A. Standard

B. Enterprise

C. Business Critical

D. Virtual Private Snowflake (VPC)

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Enterprise edition

upvoted 1 times

**Rameez1** 2 years, 4 months ago

B. Enterprise

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Enterprise

upvoted 2 times

**Avi0407** 3 years ago

Enterprise edition

upvoted 1 times

**Marge23** 3 years, 1 month ago

correct

upvoted 1 times

**happykeieiekek** 3 years, 3 months ago

https://docs.snowflake.com/en/user-guide/security-column-ddm

upvoted 1 times

**c0d3g** 3 years, 5 months ago

B. Enterprise

upvoted 1 times

**Bhagyeshbhut** 3 years, 6 months ago

Dynamic data masking require Enterprise edition or higher edition

upvoted 1 times

**rneti** 3 years, 7 months ago

The given answer is wrong. It should be (B) Enterprise Edition as per: https://docs.snowflake.com/en/user-guide/security-column-ddm

upvoted 2 times

**KarBiswa** 3 years, 8 months ago

staring is enterprise

upvoted 1 times

**mlaus** 3 years, 9 months ago

B is right

upvoted 1 times

**miadiesel** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/security-column-ddm-use.html

upvoted 4 times

**aalonso** 3 years, 10 months ago

Enterprise

upvoted 4 times

**JimXu** 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/intro-editions.html

upvoted 2 times

**tejarun** 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/intro-editions.html

upvoted 3 times

Question #215 *Topic 1*

Which services does the Snowflake Cloud Services layer manage? (Choose two.)

A. Compute resources

B. Query execution

C. Authentication

D. Data storage

E. Metadata

**Suggested Answer:** *C, E*

*Community vote distribution*

CE (100%)

**hotflake** Most Recent 2 years, 2 months ago

A. Compute resources - Compute Layer B. Query execution - Compute Layer C. Authentication - Service Layer D. Data storage - Storage Layer E. Metadata - Service Layer

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

C and E are correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

CE are correct

upvoted 1 times

**Marge23** 3 years, 1 month ago

correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**c0d3g** 3 years, 5 months ago

C. Authentication E. Metadata

upvoted 1 times

**muks86** 3 years, 8 months ago

Correct ans

upvoted 1 times

**miadiesel** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/intro-key-concepts.html

upvoted 3 times

Question #216 *Topic 1*

A company needs to allow some users to see Personally Identifiable Information (PII) while limiting other users from seeing the full value of the PII.

Which Snowflake feature will support this?

A. Row access policies

B. Data masking policies

C. Data encryption

D. Role based access control

**Suggested Answer:** *B*

*Community vote distribution*

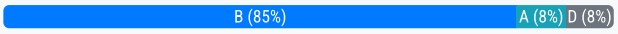


**YokoSumiGaeshi** Highly Voted 3 years, 8 months ago

If you have a table with a column including PII, masking rows will not solve the issue. What we need is to make the data in this column visible to some, and masked to some. Thus we need to use dynamic data masking.

upvoted 5 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Data Masking Policies

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Data masking policies

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

B is correct

upvoted 1 times

**c0d3g** 3 years, 5 months ago

Selected Answer: D

https://community.snowflake.com/s/article/How-to-Secure-PII-Data-with-Data-Masking

upvoted 1 times

**c0d3g** 3 years, 5 months ago

https://community.snowflake.com/s/article/How-to-Secure-PII-Data-with-Data-MaskingThe correct answer is B, Dynamic Data Masking

upvoted 1 times

**c0d3g** 3 years, 5 months ago

The correct answer is B, Dynamic Data Masking

upvoted 1 times

**DK_1** 3 years, 8 months ago

Row access policy limits the access to the row itself while Masking policy will allow to see the data value but not the actual value.

upvoted 4 times

**N89** 3 years, 8 months ago

Row acess policies can have specific roles included in the definition. The correct option is A.

upvoted 1 times

**AS314** 3 years, 10 months ago

https://community.snowflake.com/s/article/Methods-for-Securing-PII-Data-in-Snowflake I think both first and second options are correct

upvoted 3 times

Question #217 *Topic 1*

A user has unloaded data from a Snowflake table to an external stage.

Which command can be used to verify if data has been uploaded to the external stage named my_stage?

A. view @my_stage

B. list @my_stage

C. show @my_stage

D. display @my_stage

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago
list @my_stage
upvoted 1 times

**aws_god** 2 years, 7 months ago
https://docs.snowflake.com/en/sql-reference/sql/list
upvoted 1 times

**_yyukta** 2 years, 8 months ago
B. list @my_stage
upvoted 1 times

**Marge23** 3 years, 1 month ago
correct
upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago
correct
upvoted 1 times

**c0d3g** 3 years, 5 months ago
B. list @my_stage
upvoted 1 times

**EmiB** 3 years, 7 months ago
B. list @my_stage
upvoted 1 times

**Lavpak** 3 years, 11 months ago
Answer is correct List
upvoted 4 times

Question #218 *Topic 1*

Which tasks are performed in the Snowflake Cloud Services layer? (Choose two.)

A. Management of metadata

B. Computing the data

C. Maintaining Availability Zones

D. Infrastructure security

E. Parsing and optimizing queries

**Suggested Answer:** *A, E*

*Community vote distribution*

AE (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A and E are correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

AE are correct

upvoted 1 times

**Marge23** 3 years, 1 month ago

correct

upvoted 1 times

**c0d3g** 3 years, 5 months ago

A. Management of metadata E. Parsing and optimizing queries

upvoted 3 times

**girgir** 3 years, 8 months ago

why not C also for Security? https://docs.snowflake.com/en/user-guide/intro-key-concepts.html

upvoted 1 times

**girgir** 3 years, 8 months ago

why not C also for Security? https://docs.snowflake.com/en/user-guide/intro-key-concepts.htmll mean why not DThe option is "Infrastructure security" and not account security (authentication). Infrastructure is a broad term handled by Cloud Provider.

upvoted 1 times

**girgir** 3 years, 8 months ago

I mean why not DThe option is "Infrastructure security" and not account security (authentication). Infrastructure is a broad term handled by Cloud Provider.

upvoted 1 times

**girgir** 3 years, 8 months ago

I mean why not D

upvoted 1 times

**c0d3g** Highly Voted 3 years, 5 months ago

The option is "Infrastructure security" and not account security (authentication). Infrastructure is a broad term handled by Cloud Provider.

upvoted 5 times

**claudio66** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/intro-key-concepts.html

upvoted 3 times

**miadiesel** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/intro-key-concepts.html#cloud-services

upvoted 2 times

**Lavpak** 3 years, 11 months ago

Answer is correct

upvoted 2 times

Question #219 *Topic 1*

What is true about sharing data in Snowflake? (Choose two.)

A. The Data Consumer pays for data storage as well as for data computing.

B. The shared data is copied into the Data Consumer account, so the Consumer can modify it without impacting the base data of the Provider.

C. A Snowflake account can both provide and consume shared data.

D. The Provider is charged for compute resources used by the Data Consumer to query the shared data.

E. The Data Consumer pays only for compute resources to query the shared data.

**Suggested Answer:** *C, E*

*Community vote distribution*

CE (100%)

**rneti** Highly Voted 3 years, 7 months ago

The answer key (B and E) given is again wrong per the below Snowflake user-guide. The correct answer should be C and E as per popular vote.

With Secure Data Sharing, no actual data is copied or transferred between accounts. All sharing uses Snowflake's services layer and metadata store. Shared data does not take up any storage in a consumer account and therefore does not contribute to the consumer's monthly data storage charges. The only charges to consumers are for the compute resources (i.e. virtual warehouses) used to query the shared data.
(https://docs.snowflake.com/en/user-guide/data-sharing-intro.html)

upvoted 6 times

**rneti** Highly Voted 3 years, 7 months ago

The answer key (B and E) given is again wrong per the below Snowflake user-guide. The correct answer should be C and E as per popular vote. With Secure Data Sharing, no actual data is copied or transferred between accounts. All sharing uses Snowflake's services layer and metadata store. Shared data does not take up any storage in a consumer account and therefore does not contribute to the consumer's monthly data storage charges. The only charges to consumers are for the compute resources (i.e. virtual warehouses) used to query the shared data.
(https://docs.snowflake.com/en/user-guide/data-sharing-intro.html)Admins need to do something about so many objectively wrong answers. It's a bit frustrating especially when paying for these courses.

upvoted 7 times

**fellpond** 2 years, 8 months ago

Admins need to do something about so many objectively wrong answers. It's a bit frustrating especially when paying for these courses.

upvoted 5 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

C and E are correct

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct answer:CE

upvoted 1 times

**_yyukta** 2 years, 8 months ago

CE are correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

https://docs.snowflake.com/en/user-guide/data-sharing-intro

upvoted 1 times

**bonniezl** 3 years, 5 months ago

CE is correct

upvoted 1 times

**c0d3g** 3 years, 5 months ago

C. A Snowflake account can both provide and consume shared data. E. The Data Consumer pays only for compute resources to query the shared data.

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

A Consumer can be a provider for other sharing

upvoted 2 times

**aalonso** 3 years, 10 months ago

Answer is C and E

upvoted 3 times

**db1290** 3 years, 11 months ago

Answer is C and E

upvoted 2 times

**jelledv** 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/data-sharing-intro.html With Secure Data Sharing, no actual data is copied or transferred between accounts. The only charges to consumers are for the compute resources (i.e. virtual warehouses) used to query the shared data Any full Snowflake account can both provide and consume shared data

upvoted 3 times

Question #220 *Topic 1*

The following JSON is stored in a VARIANT column called src of the CAR_SALES table:

A user needs to extract the dealership information from the JSON.

How can this be accomplished?

```
{
  "customer": [
    {
       "address": San Francisco, CA",
       "name": "Jane Doe"
    }
   ],
  "date":  "2022-01-28",
  "dealership": "Town Auto Sales",
  "salesperson": {
     "id": "55"
    }
}
```

A. select src:dealership from car_sales;

B. select src.dealership from car_sales;

C. select src:Dealership from car_sales;

D. select dealership from car_sales;

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

select src:dealership from car_sales;

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. select src:dealership from car_sales;

upvoted 1 times

**Marge23** 3 years, 1 month ago

correct

upvoted 1 times

**c0d3g** 3 years, 5 months ago

Use dot notation to traverse a path in a JSON object: : . . . Optionally enclose element names in double quotes: :" "." "." ".

upvoted 1 times

**SV1122** 3 years, 10 months ago

Insert a colon : between the VARIANT column name and any first-level element: : . https://docs.snowflake.com/en/user-guide/querying-semistructured.html

upvoted 3 times

**halol** 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/querying-semistructured.html

upvoted 2 times

Question #221 *Topic 1*

Which of the following significantly improves the performance of selective point lookup queries on a table?

A. Clustering

B. Materialized Views

C. Zero-copy Cloning

D. Search Optimization Service

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**tejarun** Highly Voted 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/search-optimization-service.html#understanding-the-search-optimization-service

upvoted 6 times

**sam_art** Most Recent 2 years, 3 months ago

who says the answer is A?

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

D is correct answer

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. Search Optimization Service

upvoted 1 times

**PawPetrol** 3 years, 2 months ago

Answer D: The search optimization service can significantly improve the performance of certain types of lookup and analytical queries that use an extensive set of predicates for filtering.

upvoted 2 times

**Deeeee** 3 years, 3 months ago

Search Optimization Service

upvoted 1 times

**c0d3g** 3 years, 5 months ago

Search Optimization Service

upvoted 1 times

**AZ989** 3 years, 7 months ago

https://community.snowflake.com/s/article/Search-Optimization-When-How-To-Use

upvoted 1 times

**BigDataBB** 3 years, 9 months ago

The search optimization service aims to significantly improve the performance of certain types of queries on tables, including: Selective point lookup queries on tables. A point lookup query returns only one or a small number of distinct rows. Use case examples include: Business users who need fast response times for critical dashboards with highly selective filters. Data scientists who are exploring large data volumes and looking for specific subsets of data. Data applications retrieving a small set of results based on an extensive set of filtering predicates.

upvoted 3 times

**iamPhD** 3 years, 10 months ago

D is correct https://docs.snowflake.com/en/user-guide/search-optimization-service.html#understanding-the-search-optimization-service

upvoted 3 times

**db1290** 3 years, 11 months ago

D is correct

upvoted 3 times

**Tapasgup007** 3 years, 11 months ago

D is the Right answer

upvoted 3 times

Question #222 *Topic 1*

Which of the following accurately describes shares?

A. Tables, secure views, and secure UDFs can be shared

B. Shares can be shared

C. Data consumers can clone a new table from a share

D. Access to a share cannot be revoked once granted

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**SV1122** Highly Voted 3 years, 10 months ago

Snowflake database objects can be shared: Tables External tables Secure views Secure materialized views Secure UDFs

upvoted 6 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A is the correct option

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. Tables, secure views, and secure UDFs can be shared

upvoted 1 times

**Marge23** 3 years, 1 month ago

correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Kelvin_Yu** 3 years, 6 months ago

https://docs.snowflake.com/en/user-guide/data-sharing-intro

upvoted 1 times

**Iloveyou123** 3 years, 8 months ago

B is right?

upvoted 1 times

**Lavpak** 3 years, 11 months ago

Snowflake enables sharing of data through named Snowflake objects called shares which supports data sharing by sharing tables, secure views, and secure UDFs in our Snowflake database (Data Provider) with other Snowflake accounts (Data Consumer).

upvoted 2 times

Question #223 *Topic 1*

What are best practice recommendations for using the ACCOUNTADMIN system-defined role in Snowflake? (Choose two.)

A. Ensure all ACCOUNTADMIN roles use Multi-factor Authentication (MFA).

B. All users granted ACCOUNTADMIN role must be owned by the ACCOUNTADMIN role.

C. The ACCOUNTADMIN role must be granted to only one user.

D. Assign the ACCOUNTADMIN role to at least two users, but as few as possible.

E. All users granted ACCOUNTADMIN role must also be granted SECURITYADMIN role.

**Suggested Answer:** *A, D*

*Community vote distribution*

AD (100%)

**Lavpak** Highly Voted 3 years, 5 months ago

Answers are correct - https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html#:~:text=The%20account%20administrator%20%28i.e%20users%20with%20the%20ACCOUNTADMIN,data%2C%20and%20can%20stop%20any%20running%20SQL%20statements.

upvoted 7 times

**_yyukta** Most Recent 2 years, 2 months ago

AD are correct

upvoted 1 times

**azlovemovie** 2 years, 4 months ago

Answers are correct.

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 2 times

**Lavpak** 3 years, 5 months ago

All users assigned the ACCOUNTADMIN role should also be required to use multi-factor authentication (MFA) for login (for details, see Configuring Access Control). Assign this role to at least two users.

upvoted 4 times

Question #224 *Topic 1*

In the query profiler view for a query, which components represent areas that can be used to help optimize query performance? (Choose two.)

A. Bytes scanned

B. Bytes sent over the network

C. Number of partitions scanned

D. Percentage scanned from cache

E. External bytes scanned

**Suggested Answer:** *A, C*

*Community vote distribution*




**BigDataBB** Highly Voted 3 years, 10 months ago

I think that the right answer si C,D. Because the number of bytes scanned has no values if don't know if your query has or not make a prune, reducing the number of partition tha has scanned. I have a good query if my filters reduce the number of (micro) partiotion that the query must access to get the data.

upvoted 20 times

**bibisa** Highly Voted 3 years, 7 months ago

A. Bytes scanned represents the amount of data scanned by the query. By minimizing the number of bytes scanned, query performance can be improved. C. Number of partitions scanned represents the number of partitions read during query execution. By minimizing the number of partitions scanned, query performance can be improved. B, D, and E are incorrect statements. Bytes sent over the network, percentage scanned from cache, and external bytes scanned may provide useful information for monitoring or troubleshooting, but they are not directly related to optimizing query performance.

upvoted 9 times

**JRayan** Most Recent 1 year, 9 months ago

The components in the Snowflake query profiler that help optimize query performance are Bytes scanned (A) and Number of partitions scanned (C). Bytes scanned (A) directly impacts performance, as reducing the volume of data processed through efficient filtering or clustering improves speed and cost. Number of partitions scanned (C) reflects partition pruning efficiency. Scanning fewer partitions (via proper clustering or filters) minimizes unnecessary data access. These metrics are critical for identifying bottlenecks in data retrieval and optimizing query execution.

upvoted 2 times

**someshVashistha** 1 year, 10 months ago

Explanation: A. Bytes scanned Represents the total amount of data scanned during query execution. High values indicate that the query is scanning a large dataset, which can impact performance. Optimization: Reduce bytes scanned by applying filters, leveraging clustering, or optimizing table design. C. Number of partitions scanned Represents the number of partitions read during query execution. Scanning too many partitions can slow down query performance. Optimization: Use clustering, partition pruning, or query predicates to minimize the number of partitions scanned.

upvoted 2 times

**Kuldeepkd** 1 year, 11 months ago

A and C should be the correct answers

upvoted 1 times

**Jarturo** 2 years, 3 months ago

D & E D: Inefficient Pruning https://docs.snowflake.com/en/user-guide/ui-query-profile.html#inefficient-pruning E: Queries Too Large to Fit in Memory: https://docs.snowflake.com/en/user-guide/ui-query-profile.html#queries-too-large-to-fit-in-memory

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

C and D are correct

upvoted 2 times

**PrashantGupta1616** 2 years, 5 months ago

AC is the answer

upvoted 2 times

**Lematthew31** 2 years, 6 months ago

A. Bytes scanned: This indicates the amount of data scanned from the tables or files involved in the query. By minimizing the number of bytes scanned, you can optimize query performance by reducing the amount of data processed. C. Number of partitions scanned: This represents the number of partitions accessed during the query execution. Optimizing the number of partitions scanned can improve query performance, as accessing fewer partitions typically requires less processing time. So, the correct options are A. Bytes scanned and C. Number of partitions scanned

upvoted 2 times

**_yyukta** 2 years, 8 months ago

AC are correct

upvoted 3 times

**sfdumps** 2 years, 10 months ago

QUERY RESULT CACHE IS NOT HELPFULL

upvoted 3 times

**niho1993** 3 years, 1 month ago

reference to this question https://community.snowflake.com/s/question/0D7Do000000cgIDKAQ/detail

upvoted 2 times

**Marge23** 3 years, 1 month ago

CORRECT

upvoted 2 times

**Marge23** 3 years, 1 month ago

bytes scanned, partition scanned

upvoted 2 times

**MultiCloudIronMan** 3 years, 3 months ago

Should be correct

upvoted 1 times

**bonniezl** 3 years, 5 months ago

CD is correct

upvoted 1 times

**dmitriypo** 3 years, 5 months ago

I would go for AC. The query should not rely on cache.

upvoted 1 times

Question #225 *Topic 1*

What is the minimum Snowflake edition required for row level security?

A. Standard

B. Enterprise

C. Business Critical

D. Virtual Private Snowflake

**Suggested Answer:** *B*

*Community vote distribution*

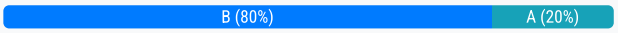


**Mallikharjuna452** Most Recent 2 years, 3 months ago

Enterprise

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Enterprise

upvoted 1 times

**SushY18** 3 years, 6 months ago

https://docs.snowflake.com/en/user-guide/security-row-intro#:~:text=Snowflake%20supports%20row-level%20security%20through%20the%20use%20of,determine%20access%20to%20rows%20in%20the%20query%20result.

upvoted 1 times

**Raju039** 3 years, 7 months ago

ENTERPRISE IS CORRECT

upvoted 1 times

**BigDataBB** 3 years, 9 months ago

Row access policies implement row-level security - https://docs.snowflake.com/en/user-guide/security-row.html

upvoted 2 times

**Chandra720** 3 years, 9 months ago

B , i.e. Minimum enterprise edition

upvoted 1 times

**Mubin_Sikilkar** 3 years, 9 months ago

Row-level security is for Enterprise or higher version

upvoted 1 times

**miadiesel** 3 years, 10 months ago

Check on the "Row Access Policies" https://docs.snowflake.com/en/user-guide/intro-editions.html

upvoted 3 times

**arpit_dataguy** 3 years, 10 months ago

Answer -- B

upvoted 2 times

**Kate0204** 3 years, 11 months ago

Selected Answer: A

correct

upvoted 2 times

**Kate0204** 3 years, 11 months ago

correctAnswer is B. Enterprise is correct

upvoted 2 times

**Raju039** 3 years, 7 months ago

Answer is B. Enterprise is correct

upvoted 1 times

**Lavpak** 3 years, 11 months ago

Answer is correct https://docs.snowflake.com/en/user-guide/intro-editions.html

upvoted 2 times

Question #226 *Topic 1*

The is the minimum Fail-safe retention time period for transient tables?

A. 1 day

B. 7 days

C. 12 hours

D. 0 days

**Suggested Answer:** *D*

*Community vote distribution*




**Mallikharjuna452** Most Recent 2 years, 3 months ago

D. 0 Days

upvoted 1 times

**AnireddySaikiranReddy** 2 years, 8 months ago

Correct Answer:D

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. 0 days

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Omkarshingare** 3 years, 7 months ago

Table Type Time Travel Retention Period (Days) Fail-safe Period (Days) Permanent 0 or 1 (for Snowflake Standard Edition) 7 0 to 90 (for Snowflake Enterprise Edition) 7 Transient 0 or 1 0 Temporary 0 or 1 0

upvoted 1 times

**Raju039** 3 years, 7 months ago

No fail safe period for Transient table

upvoted 1 times

**psilvabr** 3 years, 8 months ago

https://docs.snowflake.com/en/user-guide/tables-temp-transient "Transient tables are similar to permanent tables with the key difference that they do not have a Fail-safe period."

upvoted 2 times

**Asmit007** 3 years, 8 months ago

7days .and it is non configurable .

upvoted 1 times

**SV1122** 3 years, 10 months ago

0 days for both transient and temporary tablestemporary is 1 dayfor Time Travel

upvoted 2 times

**SV1122** 3 years, 10 months ago

Selected Answer: D

0 days for both transient and temporary tables

upvoted 2 times

**BDaly** 3 years, 10 months ago

temporary is 1 day

upvoted 1 times

**BDaly** 3 years, 10 months ago

temporary is 1 dayfor Time Travel

upvoted 1 times

**sakis213** 3 years, 10 months ago

for Time Travel

upvoted 1 times

Question #227 *Topic 1*

What is a machine learning and data science partner within the Snowflake Partner Ecosystem?

A. Informatica

B. Power BI

C. Adobe

D. Data Robot

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Data Robot

upvoted 1 times

**shankar45** 2 years, 3 months ago

Data Robot

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. Data Robot

upvoted 1 times

**Marge23** 3 years, 1 month ago

data robot

upvoted 1 times

**c0d3g** 3 years, 5 months ago

DataRobot

upvoted 1 times

**miadiesel** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/ecosystem-analytics.html

upvoted 1 times

**Lavpak** 3 years, 11 months ago

Answer is correct

upvoted 1 times

Question #228 *Topic 1*

Which statements are correct concerning the leveraging of third-party data from the Snowflake Data Marketplace? (Choose two.)

A. Data is live, ready-to-query, and can be personalized.

B. Data needs to be loaded into a cloud provider as a consumer account.

C. Data is not available for copying or moving to an individual Snowflake account.

D. Data is available without copying or moving.

E. Data transformations are required when combining Data Marketplace datasets with existing data in Snowflake.

**Suggested Answer:** *A, D*

*Community vote distribution*

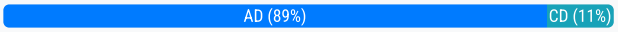


**bibisa** 3 years, 7 months ago

A. Data in the Snowflake Data Marketplace is already formatted and ready to query, and can be personalized for specific business needs. D. Data from the Snowflake Data Marketplace is accessed through Snowflake's Secure Data Sharing technology, which allows users to access the data without copying or moving it to their own account. B and C are incorrect statements. Loading data into a cloud provider as a consumer account is not required to leverage data from the Snowflake Data Marketplace, and the data can be accessed and used in a Snowflake account without restriction. E is also an incorrect statement. Data transformations may not be required when combining Data Marketplace datasets with existing data in Snowflake, as it depends on the specific data being used and how it needs to be combined or analyzed.

upvoted 5 times

**tejarun** Highly Voted 4 months ago

Selected Answer: C, D

More changes to base table increases credit consumption for materialized view maintenance

upvoted 13 times

**Mallikharjuna452** 2 years, 3 months ago

A and D are correct

upvoted 1 times

**Lavpak** Highly Voted 4 months ago

Selected Answer: C

I agree with C & D. https://docs.snowflake.com/en/user-guide/views-materialized.html#effects-of-changes-to-base-tables-on-materialized-views Answer-c --> When a base table changes, all materialized views defined on the table are updated by a background service that uses compute resources provided by Snowflake.

upvoted 6 times

**_yyukta** 2 years, 8 months ago

AD are correct

upvoted 1 times

**isolated_sniper** 4 months ago

C and E Materialized views impact your costs for both storage and compute resources: Storage: Each materialized view stores query results, which adds to the monthly storage usage for your account. Compute resources: In order to prevent materialized views from becoming out-of-date, Snowflake performs automatic background maintenance of materialized views. When a base table changes, all materialized views defined on the table are updated by a background service that uses compute resources provided by Snowflake.

upvoted 4 times

**Sk_3** 3 years ago

Why A data is no where live in mkt place it should be c

upvoted 2 times

**SV1122** 4 months ago

Selected Answer: C, D

In general, the costs are proportional to: The number of materialized views created on each base table, and the amount of data that changes in each of those materialized views when the base table changes. Any changes to micro-partitions in the base table require eventual materialized view maintenance, whether those changes are due to reclustering or DML statements run on the base table. The number of those materialized views that are clustered. Maintaining clustering (of either a table or a materialized view) adds costs. https://docs.snowflake.com/en/user-guide/views-materialized.html#effects-of-changes-to-base-tables-on-materialized-views

upvoted 4 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**titihab** Most Recent 4 months ago

I found this in Snowflake Doc : Materialized Views and Clustering Defining a clustering key on a materialized view is supported and can increase performance in many situations. However, it also adds costs. https://docs.snowflake.com/en/user-guide/views-materialized#label-materialized-views-maintenance-billing

upvoted 3 times

**Needium** 3 years, 6 months ago

A and D are the correct answers. C could be confusing but you can actually move the data to your own account if you so wish.

upvoted 1 times

**Makabaka** 3 years, 7 months ago

https://docs.snowflake.com/en/user-guide/data-share-consumers

upvoted 1 times

**AshNj** 3 years, 9 months ago

A and D

upvoted 1 times

Question #229 *Topic 1*

What impacts the credit consumption of maintaining a materialized view? (Choose two.)

A. Whether or not it is also a secure view

B. How often the underlying base table is queried

C. How often the base table changes

D. Whether the materialized view has a cluster key defined

E. How often the materialized view is queried

**Suggested Answer:** *C, D*

*Community vote distribution*

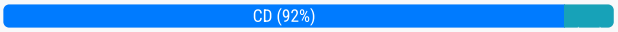


**tejarun** Highly Voted 3 years, 11 months ago

More changes to base table increases credit consumption for materialized view maintenanceAgree it's C & DI agree with C & D. https://docs.snowflake.com/en/user-guide/views-materialized.html#effects-of-changes-to-base-tables-on-materialized-views Answer-c --> When a base table changes, all materialized views defined on the table are updated by a background service that uses compute resources provided by Snowflake.

upvoted 13 times

**tejarun** Highly Voted 3 years, 11 months ago

Selected Answer: CD

More changes to base table increases credit consumption for materialized view maintenance

upvoted 13 times

**JRayan** 3 years, 8 months ago

Agree it's C & D

upvoted 1 times

**Lavpak** 3 years, 11 months ago

I agree with C & D. https://docs.snowflake.com/en/user-guide/views-materialized.html#effects-of-changes-to-base-tables-on-materialized-views Answer-c --> When a base table changes, all materialized views defined on the table are updated by a background service that uses compute resources provided by Snowflake.

upvoted 6 times

**visionary4ever** Most Recent 1 year, 2 months ago

Maintenance cost → Every time the base table changes, Snowflake automatically refreshes the MV in the background. More changes = higher cost. Query cost → When you query the MV, it uses a virtual warehouse like any table. More queries = more credits.

upvoted 1 times

**damaldon** 2 years, 2 months ago

And. CE

upvoted 1 times

**_yyukta** 2 years, 8 months ago

CD are correct

upvoted 1 times

**Sk_3** 3 years ago

CE make more sense .howdoes clustering key play role in pre computed query

upvoted 2 times

**Sk_3** 3 years ago

C and E Materialized views impact your costs for both storage and compute resources: Storage: Each materialized view stores query results, which adds to the monthly storage usage for your account. Compute resources: In order to prevent materialized views from becoming out-of-date, Snowflake performs automatic background maintenance of materialized views. When a base table changes, all materialized views defined on the table are updated by a background service that uses compute resources provided by Snowflake. This one make more sense

upvoted 2 times

**Sk_3** 3 years ago

CE make more sense .howdoes clustering key play role in pre computed queryI found this in Snowflake Doc : Materialized Views and Clustering Defining a clustering key on a materialized view is supported and can increase performance in many situations. However, it also adds costs. https://docs.snowflake.com/en/user-guide/views-materialized#label-materialized-views-maintenance-billing

upvoted 2 times

**titihab** 3 years ago

I found this in Snowflake Doc : Materialized Views and Clustering Defining a clustering key on a materialized view is supported and can increase performance in many situations. However, it also adds costs. https://docs.snowflake.com/en/user-guide/views-materialized#label-materialized-views-maintenance-billing

upvoted 4 times

**Marge23** 3 years, 1 month ago

CORRECT

upvoted 1 times

**otsumy** 3 years, 2 months ago

Why not E?

upvoted 1 times

**singhks** 3 years, 2 months ago

C and D (https://www.snowflake.com/blog/using-materialized-views-to-solve-multi-clustering-performance-problems/)

upvoted 1 times

**otsumy** 3 years, 2 months ago

Why not E?Same question here. Compute Credit should be consume, right?The question is about maintaining the view. Querying is not maintaining.

upvoted 1 times

**stornati** 3 years, 2 months ago

Same question here. Compute Credit should be consume, right?

upvoted 1 times

**stornati** 3 years, 2 months ago

Same question here. Compute Credit should be consume, right?The question is about maintaining the view. Querying is not maintaining.

upvoted 1 times

**vrfidman** 3 years ago

The question is about maintaining the view. Querying is not maintaining.

upvoted 2 times

**MultiCloudIronMan** 3 years, 3 months ago

This because al tables are clustered on Snowflake, there is no mention of re-clustering.

upvoted 1 times

**dmitriypo** 3 years, 5 months ago

Agree with C and D: https://docs.snowflake.com/en/user-guide/views-materialized#label-materialized-views-maintenance-billing

upvoted 1 times

**isolated_sniper** 3 years, 6 months ago

C and E Materialized views impact your costs for both storage and compute resources: Storage: Each materialized view stores query results, which adds to the monthly storage usage for your account. Compute resources: In order to prevent materialized views from becoming out-of-date, Snowflake performs automatic background maintenance of materialized views. When a base table changes, all materialized views defined on the table are updated by a background service that uses compute resources provided by Snowflake.

upvoted 4 times

**rneti** 3 years, 7 months ago

I would rather go with the highly voted answers (C&D) and not the key provided (D&E) based on the Snowflake Guide: https://docs.snowflake.com/en/user-guide/views-materialized.html#effects-of-changes-to-base-tables-on-materialized-views

upvoted 1 times

**psilvabr** 3 years, 8 months ago

It is C and D for this.

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

clustering and base table modifications ratio

upvoted 1 times

**mlaus** 3 years, 9 months ago

It's C&D https://docs.snowflake.com/en/user-guide/views-materialized.html#label-materialized-views-maintenance-billing

upvoted 2 times

**SV1122** 3 years, 10 months ago

In general, the costs are proportional to: The number of materialized views created on each base table, and the amount of data that changes in each of those materialized views when the base table changes. Any changes to micro-partitions in the base table require eventual materialized view maintenance, whether those changes are due to reclustering or DML statements run on the base table. The number of those materialized views that are clustered. Maintaining clustering (of either a table or a materialized view) adds costs. https://docs.snowflake.com/en/user-guide/views-materialized.html#effects-of-changes-to-base-tables-on-materialized-views

upvoted 4 times

Question #230 *Topic 1*

What COPY INTO SQL command should be used to unload data into multiple files?

A. SINGLE=TRUE

B. MULTIPLE=TRUE

C. MULTIPLE=FALSE

D. SINGLE=FALSE

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**tejarun** Highly Voted 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/data-unload-overview.html#bulk-unloading-into-single-or-multiple-files

upvoted 12 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

SINGLE = FALSE

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. SINGLE=FALSE

upvoted 1 times

**nonamae** 3 years ago

Bulk unloading, the default SINGLE= FALSE which means you can unload into multiple files

upvoted 1 times

**Marge23** 3 years, 1 month ago

CORRECT

upvoted 1 times

**PawPetrol** 3 years, 2 months ago

The default is SINGLE = FALSE (i.e. unload into multiple files) Answer: D

upvoted 1 times

**srj_mehta** 3 years, 4 months ago

ANSWER IS D The COPY INTO command provides a copy option (SINGLE) for unloading data into a single file or multiple files. The default is SINGLE = FALSE (i.e. unload into multiple files).

upvoted 1 times

**EmiB** 3 years, 7 months ago

D. SINGLE=FALSE

upvoted 1 times

**psilvabr** 3 years, 8 months ago

D SINGLE=FALSE

upvoted 1 times

**Varma_Saraswathula** 3 years, 8 months ago

SINGLE = FALSE

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

single is false

upvoted 1 times

**vinnno** 3 years, 11 months ago

The default is SINGLE = FALSE (i.e. unload into multiple files).

upvoted 4 times

Question #231 *Topic 1*

When cloning a database containing stored procedures and regular views, that have fully qualified table references, which of the following will occur?

A. The cloned views and the stored procedures will reference the cloned tables in the cloned database.

B. An error will occur, as views with qualified references cannot be cloned.

C. An error will occur, as stored objects cannot be cloned.

D. The stored procedures and views will refer to tables in the source database.

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**Rajivnb** Highly Voted 2 years, 6 months ago

Fully Qualified table reference means, specifying the table name like the below, . . . The questions says, the DB containing stored procedures and regular views (reason it was told a regular view is, for secure view the table cross reference across DB needs a REFERENCE GRANT) that have a fully qualified table references. If its a not a qualified name, then the Table will be looked into the current DB, which is the DB in which the cloned SP and Views are created.

upvoted 7 times

**halol** Highly Voted 3 years, 5 months ago

https://docs.snowflake.com/en/sql-reference/sql/create-clone.html

upvoted 5 times

**_yyukta** Most Recent 2 years, 2 months ago

D. The stored procedures and views will refer to tables in the source database

upvoted 2 times

**dmitriypo** 2 years, 12 months ago

For sure D. Just tested this.

upvoted 2 times

**psilvabr** 3 years, 1 month ago

Not A for sure, D

upvoted 1 times

Question #232 *Topic 1*

When loading data into Snowflake, how should the data be organized?

A. Into single files with 100-250 MB of compressed data per file

B. Into single files with 1-100 MB of compressed data per file

C. Into files of maximum size of 1 GB of compressed data per file

D. Into files of maximum size of 4 GB of compressed data per file

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A is correct answer

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. Into single files with 100-250 MB of compressed data per file

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**c0d3g** 3 years, 5 months ago

Into single files with 100-250 MB of compressed data per file

upvoted 1 times

**Lavpak** 3 years, 11 months ago

Answer is correct

upvoted 3 times

Question #233 *Topic 1*

Which of the following objects can be directly restored using the UNDROP command? (Choose two.)

A. Schema

B. View

C. Internal stage

D. Table

E. User

F. Role

**Suggested Answer:** *A, D*

*Community vote distribution*

AD (100%)

**SV1122** Highly Voted 3 years, 10 months ago

Account Objects: UNDROP DATABASE Database Objects: UNDROP SCHEMA UNDROP TABLE UNDROP TAG

upvoted 8 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A and D are correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

AD are correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Lavpak** 3 years, 11 months ago

Answer is correct https://docs.snowflake.com/en/sql-reference/sql/undrop.html

upvoted 2 times

Question #234 *Topic 1*

Which Snowflake SQL statement would be used to determine which users and roles have access to a role called MY_ROLE?

A. SHOW GRANTS OF ROLE MY_ROLE

B. SHOW GRANTS TO ROLE MY_ROLE

C. SHOW GRANTS FOR ROLE MY_ROLE

D. SHOW GRANTS ON ROLE MY_ROLE

**Suggested Answer:** *A*

*Community vote distribution*

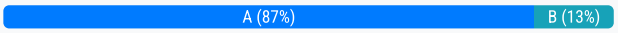


**miadiesel** Highly Voted 3 years, 10 months ago

A: SHOW GRANTS OF ROLE shows who as been granted (Correct) B: SHOW GRANTS TO ROLE show privileges granted https://docs.snowflake.com/en/sql-reference/sql/show-grants.html

upvoted 12 times

**Avi0407** Most Recent 3 years ago

A is the answer.. List all roles and users who have been granted the analyst role: SHOW GRANTS OF ROLE analyst; List all privileges granted to the analyst role: SHOW GRANTS TO ROLE analyst;

upvoted 2 times

**Avi0407** 3 years ago

List all privileges granted to the ROLE role: https://docs.snowflake.com/en/sql-reference/sql/show-grants SHOW GRANTS TO ROLE analyst; +----------+----------+----------+----------+----------+----------+----------+ | created_on | privilege | granted_on | name | granted_to | grant_option | granted_by | |----------+----------+----------+----------+----------+----------+----------+ | Wed, 17 Dec 2014 18:19:37 -0800 | CREATE WAREHOUSE | ACCOUNT | DEMOENV | ANALYST | false | SYSADMIN | +----------+----------+----------+----------+----------+----------+----------+

upvoted 3 times

**Afzy** 3 years ago

Answer is B SHOW GRANTS TO ROLE ;

upvoted 1 times

**Rajivnb** 3 years ago

OF is correct.

upvoted 1 times

**Marge23** 3 years, 1 month ago

correct

upvoted 1 times

**timcooks** 3 years, 1 month ago

SHOW GRANTS TO ROLE //Is the correct one else One can check on worksheet.

upvoted 1 times

**PawPetrol** 3 years, 2 months ago

Ans is A

upvoted 1 times

**OscarWest2023** 3 years, 5 months ago

https://docs.snowflake.com/en/sql-reference/sql/show-grants#examples

upvoted 2 times

**okurok** 3 years, 7 months ago

https://docs.snowflake.com/en/sql-reference/sql/show-grants#:~:text=to%20the%20share.-,SHOW%20GRANTS%20OF,-...

upvoted 1 times

**SexyPotato** 3 years, 8 months ago

A is correct.

upvoted 1 times

**Andywu213** 3 years, 10 months ago

Answer is A SHOW GRANTS OF... ROLE role_name Lists all users and roles to which the role has been granted.

upvoted 2 times

**SV1122** 3 years, 10 months ago

MY_ROLE doesn't have access to users and roles. It is the other way round - Users and roles were given access to MY_ROLE so SHOW GRANTS OF ROLE MY_ROLE.

upvoted 3 times

**db1290** 3 years, 11 months ago

Answer is A. SHOW OF ROLE and SHOW TO ROLE has different outputs.

upvoted 2 times

**stornati** 3 years, 11 months ago

It could A or B Syntax SHOW GRANTS SHOW GRANTS ON ACCOUNT SHOW GRANTS ON SHOW GRANTS TO { ROLE | USER | SHARE } SHOW GRANTS OF ROLE SHOW GRANTS OF SHARE SHOW FUTURE GRANTS IN SCHEMA { } SHOW FUTURE GRANTS IN DATABASE { } SHOW FUTURE GRANTS TO ROLE https://docs.snowflake.com/en/sql-reference/sql/show-grants.html

upvoted 1 times

**tejarun** 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/sql/show-grants.html

upvoted 2 times

Question #235 *Topic 1*

What is the MINIMUM edition of Snowflake that is required to use a SCIM security integration?

A. Business Critical Edition

B. Standard Edition

C. Virtual Private Snowflake (VPS)

D. Enterprise Edition

**Suggested Answer:** *B*

*Community vote distribution*

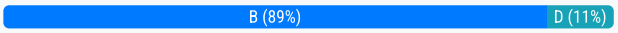


**yogen** Highly Voted 2 years, 4 months ago

https://docs.snowflake.com/en/user-guide/admin-security.html Its available for ALL editions so Standard edition is the answer

upvoted 5 times

**visionary4ever** Most Recent 1 year, 2 months ago

Enterprise unlocks identity integrations (SSO, MFA, SCIM).

upvoted 1 times

**SexyPotato** 2 years, 2 months ago

B is right.

upvoted 1 times

**fahfouhi94** 2 years, 3 months ago

I just tested, available on standard edition also

upvoted 3 times

**Babsbk** 2 years, 4 months ago

Its available for ALL editions so Standard edition is the answe

upvoted 4 times

Question #236 *Topic 1*

A user created a transient table and made several changes to it over the course of several days. Three days after the table was created, the user would like to go back to the first version of the table.

How can this be accomplished?

A. Use Time Travel, as long as DATA_RETENTION_TIME_IN_DAYS was set to at least 3 days.

B. The transient table version cannot be retrieved after 24 hours.

C. Contact Snowflake Support to have the data retrieved from Fail-safe storage.

D. Use the FAIL_SAFE parameter for Time Travel to retrieve the data from Fail-safe storage.

**Suggested Answer:** *B*

*Community vote distribution*

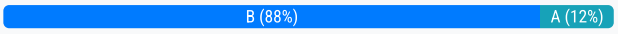


**0e504b5** Most Recent 2 years, 3 months ago

https://docs.snowflake.com/en/user-guide/tables-temp-transient

upvoted 3 times

**timcooks** 2 years, 7 months ago

If the snowflake edition is Enterprise and can go up-to to 90 days. If time travel set to 90 days for Transient tables.

upvoted 2 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 2 times

**girgir** 3 years, 2 months ago

A use time travel https://community.snowflake.com/s/article/Time-Travel-Inherited-DATA-RETENTION-TIME-IN-DAYS-Parameter-Retained-in-Transient-Tables

upvoted 3 times

**fahfouhi94** 3 years, 3 months ago

correct answer

upvoted 4 times

**SV1122** 3 years, 4 months ago

For transient table the time travel is from 0 to 1 days.

upvoted 2 times

**Lavpak** 3 years, 5 months ago

Correct Answer

upvoted 3 times

Question #237 *Topic 1*

When reviewing the load for a warehouse using the load monitoring chart, the chart indicates that a high volume of queries is always queuing in the warehouse.

According to recommended best practice, what should be done to reduce the queue volume? (Choose two.)

A. Use multi-clustered warehousing to scale out warehouse capacity.

B. Scale up the warehouse size to allow queries to execute faster.

C. Stop and start the warehouse to clear the queued queries.

D. Migrate some queries to a new warehouse to reduce load.

E. Limit user access to the warehouse so fewer queries are run against it.

**Suggested Answer:** *A, D*

*Community vote distribution*

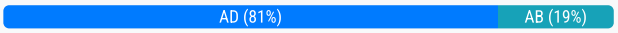


**danilchezz** Highly Voted 3 years, 11 months ago

Selected Answer: AD

This is per Snowflake documentation: https://docs.snowflake.com/en/user-guide/warehouses-load-monitoring.html#slow-query-performance

upvoted 12 times

**danilchezz** Highly Voted 3 years, 11 months ago

This is per Snowflake documentation: https://docs.snowflake.com/en/user-guide/warehouses-load-monitoring.html#slow-query-performanceIf the running query load is high or there's queuing, consider - starting a separate warehouse and moving queued queries to that warehouse. - Alternatively, if you are using multi-cluster warehouses, you could change your multi-cluster settings to add additional clusters to handle higher concurrency going forward.

upvoted 12 times

**EmiB** 3 years, 7 months ago

If the running query load is high or there's queuing, consider - starting a separate warehouse and moving queued queries to that warehouse. - Alternatively, if you are using multi-cluster warehouses, you could change your multi-cluster settings to add additional clusters to handle higher concurrency going forward.

upvoted 3 times

**PrasantSadangi** Most Recent 1 year, 7 months ago

consider starting a separate warehouse and moving queued queries to that warehouse.

upvoted 1 times

**vivekmani2021** 2 years, 3 months ago

You are trying to solve for concurrency, so either scale out or add new warehouses

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct answer:AB

upvoted 1 times

**nexerSnow** 2 years, 5 months ago

A&B r correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Needium** 3 years, 6 months ago

if queries execute faster, there will be lesser queries in the queue. moving to another virtual warehouse may be a temporary fix, but when it's a persistent problem then you should first scale out your virtual warehouse. If scaling out does not fully address the issue, you should consider scaling up as well

upvoted 2 times

**Raju039** 3 years, 7 months ago

If the running query load is high or there's queuing, consider starting a separate warehouse and moving queued queries to that warehouse. Alternatively, if you are using multi-cluster warehouses, you could change your multi-cluster settings to add additional clusters to handle higher concurrency going forward.

upvoted 3 times

**OTE** 3 years, 7 months ago

Tricky question as scaling up (B) could potentially also help. However, I think it's A/D because the question is about loading data and, thus, concurrency, not complex queries.

upvoted 4 times

**girgir** 3 years, 8 months ago

AB, not AD, the key word here is 'always' which indicates the action needed is for following computes but not necessarily current running queries. moving queued queries to new warehouse is temp solution. scale up and scale out is the permanent solution. https://docs.snowflake.com/en/user-guide/warehouses-load-monitoring.html#slow-query-performance

upvoted 3 times

**girgir** 3 years, 8 months ago

AB, not AD, the key word here is 'always' which indicates the action needed is for following computes but not necessarily current running queries. moving queued queries to new warehouse is temp solution. scale up and scale out is the permanent solution. https://docs.snowflake.com/en/user-guide/warehouses-load-monitoring.html#slow-query-performanceIt's AD. Scaling up(B) is needed only when Query performance is slow.

upvoted 3 times

**Raju039** 3 years, 7 months ago

It's AD. Scaling up(B) is needed only when Query performance is slow.

upvoted 2 times

**n21007** 3 years, 8 months ago

https://docs.snowflake.com/en/user-guide/warehouses-load-monitoring.html#slow-query-performance

upvoted 2 times

**fahfouhi94** 3 years, 9 months ago

correct answer

upvoted 1 times

**sailoo** 3 years, 9 months ago

AD for sure

upvoted 1 times

**nintendogamer64** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/warehouses-load-monitoring.html#slow-query-performance

upvoted 1 times

**nintendogamer64** 3 years, 10 months ago

AB: If the running query load is high or there's queuing, consider starting a separate warehouse and moving queued queries to that warehouse. Alternatively, if you are using multi-cluster warehouses, you could change your multi-cluster settings to add additional clusters to handle higher concurrency going forward.

upvoted 2 times

**nintendogamer64** 3 years, 10 months ago

AB: If the running query load is high or there's queuing, consider starting a separate warehouse and moving queued queries to that warehouse. Alternatively, if you are using multi-cluster warehouses, you could change your multi-cluster settings to add additional clusters to handle higher concurrency going forward.source: https://docs.snowflake.com/en/user-guide/warehouses-load-monitoring.html#slow-query-performance

upvoted 2 times

**nintendogamer64** 3 years, 10 months ago

source: https://docs.snowflake.com/en/user-guide/warehouses-load-monitoring.html#slow-query-performance

upvoted 2 times

**SV1122** 3 years, 10 months ago

This is interesting. A,B,D all 3 are correct according to https://docs.snowflake.com/en/user-guide/warehouses-load-monitoring.html#slow-query-performanceif queries execute faster, there will be lesser queries in the queue. moving to another virtual warehouse may be a temporary fix, but when it's a persistent problem then you should first scale out your virtual warehouse. If scaling out does not fully address the issue, you should consider scaling up as well

upvoted 1 times

**SV1122** 3 years, 10 months ago

This is interesting. A,B,D all 3 are correct according to https://docs.snowflake.com/en/user-guide/warehouses-load-monitoring.html#slow-query-performance

upvoted 1 times

**halol** 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/warehouses-load-monitoring.html#slow-query-performance

upvoted 2 times

Question #238 *Topic 1*

Which of the following features, associated with Continuous Data Protection (CDP), require additional Snowflake-provided data storage? (Choose two.)

A. Tri-Secret Secure

B. Time Travel

C. Fail-safe

D. Data encryption

E. External stages

**Suggested Answer:** *B, C*

*Community vote distribution*

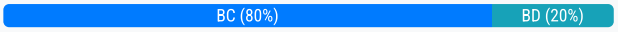


**Mallikharjuna452** 2 years, 3 months ago

B and C are Correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**c0d3g** 3 years, 5 months ago

Time Travel Fail-safe

upvoted 1 times

**chandra_157_447** 3 years, 5 months ago

correct answer is TimeTravel and Data Encryption. If you use the same link provided by other comments, you can see that Time Travel and Fail Safe as a same feature.But even if you configure Time-travel to 0 days, in case you have an Edition with Fail-Safe and the object is not transient or temporary, you'll paying for the Fail-Safe for 7 days.

upvoted 1 times

**chandra_157_447** 3 years, 5 months ago

Selected Answer: BD

correct answer is TimeTravel and Data Encryption. If you use the same link provided by other comments, you can see that Time Travel and Fail Safe as a same feature.

upvoted 1 times

**06c63d0** Most Recent 2 years, 2 months ago

But even if you configure Time-travel to 0 days, in case you have an Edition with Fail-Safe and the object is not transient or temporary, you'll paying for the Fail-Safe for 7 days.

upvoted 1 times

**Lavpak** 3 years, 11 months ago

Correct Answer https://docs.snowflake.com/en/user-guide/data-cdp.html both Time Travel and Fail-safe require additional data storage, which has associated fees

upvoted 2 times

Question #239 *Topic 1*

Where can a user find and review the failed logins of a specific user for the past 30 days?

A. The USERS view in ACCOUNT_USAGE

B. The LOGIN_HISTORY view in ACCOUNT_USAGE

C. The ACCESS_HISTORY view in ACCOUNT_USAGE

D. The SESSIONS view in ACCOUNT_USAGE

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B is correct Answer

upvoted 1 times

**fahfouhi94** 3 years, 9 months ago

correct answer

upvoted 1 times

**SV1122** 3 years, 10 months ago

https://docs.snowflake.com/en/sql-reference/account-usage/login_history.html

upvoted 1 times

Question #240 *Topic 1*

What is the purpose of an External Function?

A. To call code that executes outside of Snowflake

B. To run a function in another Snowflake database

C. To share data in Snowflake with external parties

D. To ingest data from on-premises data sources

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A is correct answer

upvoted 1 times

**KAMAL97** 2 years, 6 months ago

Even C makes sense as next priority Answer

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**fahfouhi94** 3 years, 9 months ago

https://docs.snowflake.com/en/sql-reference/external-functions.html

upvoted 2 times

**Lavpak** 3 years, 11 months ago

Correct Answer https://docs.snowflake.com/en/sql-reference/external-functions-introduction.html#what-is-an-external-function

upvoted 4 times

Question #241 *Topic 1*

Which of the following statements apply to Snowflake in terms of security? (Choose two.)

A. Snowflake leverages a Role-Based Access Control (RBAC) model.

B. Snowflake requires a user to configure an IAM user to connect to the database.

C. All data in Snowflake is encrypted.

D. Snowflake can run within a user's own Virtual Private Cloud (VPC).

E. All data in Snowflake is compressed.

**Suggested Answer:** *A, C*

*Community vote distribution*

AC (100%)

**ntechkuma** Highly Voted 3 years, 7 months ago

AC is correct answer. Moderator should rectity the answer choice

upvoted 9 times

**stornati** Highly Voted 3 years, 11 months ago

You cannot run Snowflake in your own VPC

upvoted 5 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A and C correct

upvoted 1 times

**pranalig** 2 years, 4 months ago

Correct:AC

upvoted 1 times

**azlovemovie** 2 years, 10 months ago

A and C are the right answers.

upvoted 1 times

**Afzy** 3 years ago

it is confusing, some please explain why it cam=n't be E ?I guess compression is about storage optimality ! The question here is for security

upvoted 1 times

**Afzy** 3 years ago

it is confusing, some please explain why it cam=n't be E ?

upvoted 1 times

**yrajo** 2 years, 11 months ago

I guess compression is about storage optimality ! The question here is for security

upvoted 4 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**RajuNaik** 3 years, 6 months ago

AC is correct answer.

upvoted 1 times

**psilvabr** 3 years, 8 months ago

Snowflake is a native 100% public cloud solution, you cannot host it on your OWN VPC. All data micro partitions are encrypted.

upvoted 4 times

**Josa2** 3 years, 8 months ago

What the hell? How come it's D? Who is curating this?

upvoted 1 times

**fahfouhi94** 3 years, 9 months ago

yes, AC is correct answer

upvoted 1 times

**nzpentest** 3 years, 11 months ago

A C I just finished reading the security section.

upvoted 4 times

**jelledv** 3 years, 11 months ago

You cannot run Snowflake in your own VPC. All data is encrypted.

upvoted 4 times

Question #242 *Topic 1*

A single user of a virtual warehouse has set the warehouse to auto-resume and auto-suspend after 10 minutes. The warehouse is currently suspended and the user performs the following actions:

1. Runs a query that takes 3 minutes to complete
2. Leaves for 15 minutes
3. Returns and runs a query that takes 10 seconds to complete
4. Manually suspends the warehouse as soon as the last query was completed

When the user returns, how much billable compute time will have been consumed?

A. 4 minutes

B. 10 minutes

C. 14 minutes

D. 24 minutes

**Suggested Answer:** *C*

*Community vote distribution*

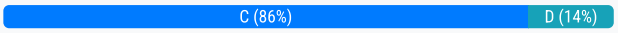


**stornati** Highly Voted 3 years, 5 months ago

I think it is 14 mins because: 3 Minutes for running first time (starting the WH and first execution) Leave for 15 minutes. WH will be iddle after 10 mins. ==> 10 + 3 New execuion = Minimal is 1 minute billed. so: 10+3+1 = 14Why after the second run WH suspends after 1 min if autosuspend set to 10 min?4th point in the question says the user suspended the Warehouse manually immediately after the second query is completed. if the user would have left the warehouse active then it would have been 24.

upvoted 16 times

**stornati** Highly Voted 3 years, 5 months ago

Selected Answer: C

I think it is 14 mins because:

3 Minutes for running first time (starting the WH and first execution) Leave for 15 minutes. WH will be iddle after 10 mins. ==> 10 + 3 New execuion = Minimal is 1 minute billed.

so: 10+3+1 = 14

upvoted 16 times

**v1tjkee** 2 years, 7 months ago

Why after the second run WH suspends after 1 min if autosuspend set to 10 min?

upvoted 1 times

**v1tjkee** 2 years, 7 months ago

Why after the second run WH suspends after 1 min if autosuspend set to 10 min?4th point in the question says the user suspended the Warehouse manually immediately after the second query is completed. if the user would have left the warehouse active then it would have been 24.

upvoted 1 times

**Rajivnb** 2 years, 6 months ago

4th point in the question says the user suspended the Warehouse manually immediately after the second query is completed. if the user would have left the warehouse active then it would have been 24.

upvoted 2 times

**PrasantSadangi** Most Recent 1 year, 7 months ago

14 Mins

upvoted 1 times

**deeshubh** 2 years ago

3 Minutes for running first time (starting the WH and first execution) Leave for 15 minutes. WH will be idle after 10 mins. ==> 10 + 3 New execution 10seconds = Minimal is 1 minute billed. total :- 14

upvoted 2 times

**_yyukta** 2 years, 2 months ago

C. 14 minutes

upvoted 1 times

**SCube** 2 years, 6 months ago

If the warehouse is currently suspended , how does the 3 min query get executed, if someone can answer plzbc auto resume is on

upvoted 1 times

**SCube** 2 years, 6 months ago

If the warehouse is currently suspended , how does the 3 min query get executed, if someone can answer plz

upvoted 1 times

**sandalz** 2 years, 3 months ago

bc auto resume is on

upvoted 2 times

**jAtlas7** 2 years, 7 months ago

First run: 3 minute query plus the default 10 minutes until the warehouse is suspended is 13 minutes. Second run: 10 second query (billed as 1 minute) and then immediate manual suspension is only still 1 minute. 13+1 = 14 minutes.

upvoted 2 times

**stornati** 2 years, 8 months ago

3 +10 (idle before auto suspend) +1 (run quey) + 10 (idle before auto suspend, since user gone without suspend)

upvoted 3 times

**v1tjkee** 2 years, 10 months ago

Selected Answer: D

I couldn't understand why after the first query (3min) we have 10 minutes before WH is suspended, then after the second query (10sec) we don't have 10 minutes before WH is suspended?

upvoted 1 times

**v1tjkee** 2 years, 10 months ago

I couldn't understand why after the first query (3min) we have 10 minutes before WH is suspended, then after the second query (10sec) we don't have 10 minutes before WH is suspended?BECAUSE MANULLY SUSPEND

upvoted 1 times

**samudu141** 2 years, 9 months ago

BECAUSE MANULLY SUSPEND

upvoted 2 times

**Faouzi5** 2 years, 10 months ago

The correct answer is 24: 3 +10 (idle before auto suspend) +1 + 10 (idle before auto suspend)

upvoted 2 times

**Faouzi5** 2 years, 10 months ago

The correct answer is 24: 3 +10 (idle before auto suspend) +1 + 10 (idle before auto suspend)Cannot be 24 because the user suspends the warehouse after the last query is complete.

upvoted 2 times

**SCube** 2 years, 6 months ago

Cannot be 24 because the user suspends the warehouse after the last query is complete.

upvoted 1 times

**RajuNaik** 3 years ago

14 min

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

3 min of the first query running + 10 min berfore the auto-suspend + 1 min (min for a new exection)

upvoted 4 times

Question #243 *Topic 1*

What can be used to view warehouse usage over time? (Choose two.)

A. The LOAD HISTORY view

B. The query history view

C. The SHOW WAREHOUSES command

D. The WAREHOUSE_METERING_HISTORY view

E. The billing and usage tab in the Snowflake web UI

**Suggested Answer:** *D, E*

*Community vote distribution*

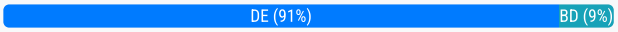


**halol** Highly Voted 3 years, 10 months ago

Maybe D&E

upvoted 10 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

D and E are correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D And E are correct

upvoted 1 times

**Afzy** 3 years ago

answer is D & E

upvoted 2 times

**singhks** 3 years, 2 months ago

A, D : https://medium.com/snowflake/snowflake-warehouse-load-and-metering-analysis-f900e5448c8a

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

correct

upvoted 1 times

**mayank** 3 years, 8 months ago

should be the correct answer

upvoted 3 times

**fahfouhi94** 3 years, 9 months ago

DE is correct answer query history view specify the wh used but no info about usage consumption

upvoted 2 times

**BigDataBB** 3 years, 9 months ago

WAREHOUSE_METERING_HISTORY View This Account Usage view can be used to return the hourly credit usage for a single warehouse (or all the warehouses in your account) within the last 365 days (1 year). https://docs.snowflake.com/en/sql-reference/account-usage/warehouse_metering_history.html Snowsight can be used to view complute cost https://docs.snowflake.com/en/user-guide/cost-exploring-compute.html

upvoted 3 times

**yogen** 3 years, 10 months ago

DE Snowsight can be used to view complute cost https://docs.snowflake.com/en/user-guide/cost-exploring-compute.html

upvoted 3 times

**SV1122** 3 years, 10 months ago

Query history https://docs.snowflake.com/en/sql-reference/account-usage.html#label-account-usage-warehouse-performance-query

upvoted 1 times

Question #244 *Topic 1*

What actions will prevent leveraging of the ResultSet cache? (Choose two.)

A. Removing a column from the query SELECT list

B. Stopping the virtual warehouse that the query is running against

C. Clustering of the data used by the query

D. Executing the RESULTS_SCAN() table function

E. Changing a column that is not in the cached query

**Suggested Answer:** *A, E*

*Community vote distribution*




**SV1122** Highly Voted 3 years, 4 months ago

Selected Answer: AC

C is explained here "The table's micro-partitions have not changed (e.g. been reclustered or consolidated) due to changes to other data in the table."

https://docs.snowflake.com/en/user-guide/querying-persisted-results.html

upvoted 9 times

**SV1122** Highly Voted 3 years, 4 months ago

C is explained here "The table's micro-partitions have not changed (e.g. been reclustered or consolidated) due to changes to other data in the table."
https://docs.snowflake.com/en/user-guide/querying-persisted-results.htmlACE are correct.

upvoted 9 times

**Mike722** 2 years, 2 months ago

ACE are correct.

upvoted 1 times

**Goransoll** Most Recent 7 months, 4 weeks ago

AE A. Removing a column from the query SELECT list Snowflake's ResultSet cache is exact-match only. If you change anything in the query text — even removing a column — Snowflake treats it as a different query and will not reuse the cached result. E. Changing a column that is not in the cached query Even if the changed column wasn't selected, Snowflake still invalidates the ResultSet cache because the underlying data used by the query has changed. ResultSet cache requires no changes to any table data referenced by the query. Why not C: Clustering affects micro-partition pruning and performance, not the ResultSet cache.

upvoted 1 times

**ChewyRemedy** 1 year, 8 months ago

If you modify the query, such as by removing or adding a column, the ResultSet cache cannot be leveraged since the query is now different from the one that was previously cached. Clustering changes the physical storage of the data, which affects the query results and prevents the use of the ResultSet cache.

upvoted 1 times

**JRayan** 1 year, 9 months ago

AE is the correct answer

upvoted 1 times

**MultiCloudIronMan** 1 year, 9 months ago

Clustering does not change data or query results but makes a query more efficient.Clustering changes the micro-partitions. Of course, here clustering means re-clustering.

upvoted 2 times

**MultiCloudIronMan** 1 year, 9 months ago

Selected Answer: AE

Clustering does not change data or query results but makes a query more efficient.

upvoted 2 times

**37ceea2** 1 year, 8 months ago

Clustering changes the micro-partitions. Of course, here clustering means re-clustering.

upvoted 1 times

**NPass** 1 year, 9 months ago

Re-clustering data does not invalidate the ResultSet cache. The cache is based on query text and the state of the data, not how the data is physically stored. Answer: AE

upvoted 3 times

**abhijit2683** 2 years ago

AE are correct

upvoted 3 times

**rogerek077** 2 years, 1 month ago

If "Changing a column that is not in the cached query" means changing such column's name or updating the data of that column then ACE are correct answers.

upvoted 1 times

**_yyukta** 2 years, 2 months ago

A and C are correct

upvoted 1 times

**majew** 2 years, 3 months ago

I think it's not REclustering so C is not correct

upvoted 2 times

**sandy91** 2 years, 3 months ago

By clustering you are reorganizing the micro partitions thereby preventing resultset cache to reuse.. So A and C ae correct

upvoted 2 times

**umidjon03** 2 years, 5 months ago

ACE are correct. I've checked the case of E, and persisted query cache was not retrieved

upvoted 2 times

**Jangelovska** 2 years, 5 months ago

A -- Correct E -- Correct Tested with update on column that is not part of the cached query, after that cached cannot be used.

upvoted 1 times

**NEMMELA** 2 years, 8 months ago

AC is correct

upvoted 1 times

**arsalanramin** 2 years, 8 months ago

Should be A, B

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**jayeshtest** 2 years, 10 months ago

Clustering will not prevent use of cached results, reclustering or consolidation will.

upvoted 4 times

Question #245 *Topic 1*

Which statement is true about running tasks in Snowflake?

A. A task can be called using a CALL statement to run a set of predefined SQL commands.

B. A task allows a user to execute a single SQL statement/command using a predefined schedule.

C. A task allows a user to execute a set of SQL commands on a predefined schedule.

D. A task can be executed using a SELECT statement to run a predefined SQL command.

**Suggested Answer:** *B*

*Community vote distribution*

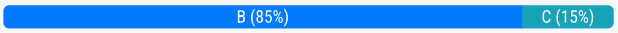


**SV1122** Highly Voted 3 years, 10 months ago

A task can execute any one of the following types of SQL code: Single SQL statement Call to a stored procedure Procedural logic using Snowflake Scripting Developer Guide https://docs.snowflake.com/en/user-guide/tasks-intro.html

upvoted 9 times

**niho1993** Most Recent 1 year, 1 month ago

A Task in Snowflake is designed to execute a single SQL statement (could be a COPY INTO, a MERGE, a CALL to a stored procedure, etc.) on a defined schedule or dependency tree. If you need to run multiple SQL statements, you must wrap them inside a stored procedure and then schedule the task to CALL it. Tasks cannot be triggered by SELECT or CALL directly; they're scheduled.

upvoted 1 times

**user_1011** 1 year, 5 months ago

https://docs.snowflake.com/en/sql-reference/sql/create-task As per the snowflake documentation, multiple SQL statments can be executed in a task through SNOWSQL by setting up a delimiter

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B is correct answer

upvoted 1 times

**YoshiSamoZou** 2 years, 4 months ago

agree b

upvoted 1 times

**deeshubh** 2 years, 6 months ago

Multiple SQL statements using SnowSQL Create task executes multiple SQL statements. Example, task modifies the TIMESTAMP_OUTPUT_FORMAT for the session and then queries the CURRENT_TIMESTAMP function. The SQL code in the task definition includes multiple statements. To execute the CREATE TASK statement, you must temporarily set a character other than a semicolon as the delimiter for SQL statements; otherwise, the CREATE TASK statement would return a user error. The command to change the SQL delimiter in SnowSQL is !set sql_delimiter = . https://docs.snowflake.com/en/sql-reference/sql/create-task Now Decide B or C, As per my understanding, may be now they have add this feature, so as per old snowflake its B but for now its C

upvoted 2 times

**Prashn** 2 years, 8 months ago

https://docs.snowflake.com/en/sql-reference/sql/create-task Multiple SQL statements using SnowSQL !set sql_delimiter=/ CREATE OR REPLACE TASK test_logging USER_TASK_MANAGED_INITIAL_WAREHOUSE_SIZE = 'XSMALL' SCHEDULE = 'USING CRON 0 * * * * America/Los_Angeles' AS BEGIN ALTER SESSION SET TIMESTAMP_OUTPUT_FORMAT = 'YYYY-MM-DD HH24:MI:SS.FF'; SELECT CURRENT_TIMESTAMP; END;/ !set sql_delimiter=";"

upvoted 2 times

**_yyukta** 2 years, 8 months ago

B. A task allows a user to execute a single SQL statement/command using a predefined schedule.

upvoted 1 times

**OscarWest2023** 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/tasks-intro

upvoted 1 times

**dmitriypo** 3 years, 5 months ago

Agree - B

upvoted 1 times

**Prajdeo** 3 years, 8 months ago

A task can execute any one of the following types of SQL code: Single SQL statement Call to a stored procedure Procedural logic using Snowflake Scripting Developer Guide https://docs.snowflake.com/en/user-guide/tasks-intro

upvoted 3 times

**snowcert** 3 years, 10 months ago

B is the answer

upvoted 2 times

**Rob__C** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/tasks-intro.htmlA task can execute any one of the following types of SQL code: Single SQL statement

upvoted 2 times

**Rob__C** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/tasks-intro.html

upvoted 2 times

**SnowProCertDec22** 3 years, 10 months ago

A task can execute any one of the following types of SQL code: Single SQL statement

upvoted 2 times

**halol** 3 years, 10 months ago

Most probably B as task can run a single SQL statement

upvoted 2 times

Question #246 *Topic 1*

Which data types does Snowflake support when querying semi-structured data? (Choose two.)

A. VARIANT

B. VARCHAR

C. XML

D. ARRAY

E. BLOB

**Suggested Answer:** *A, D*

*Community vote distribution*




**AlexbDku** Highly Voted 3 years, 10 months ago

A&D https://docs.snowflake.com/en/sql-reference/data-types-semistructured.html

upvoted 6 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Variant and array

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A and D are correct

upvoted 1 times

**singhks** 3 years, 2 months ago

A and D are correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**pab1004** 3 years, 7 months ago

A y D correct

upvoted 1 times

**miadiesel** 3 years, 10 months ago

https://docs.snowflake.com/en/sql-reference/data-types-semistructured.html

upvoted 1 times

**SnowProCertDec22** 3 years, 10 months ago

d y a are correct

upvoted 1 times

**halol** 3 years, 11 months ago

Selected Answer: AB

Variant and Array

upvoted 2 times

**halol** 3 years, 11 months ago

Variant and ArrayA,D Variant and Array

upvoted 2 times

**halol** 3 years, 11 months ago

A,D Variant and Array

upvoted 3 times

**harshagc180** 3 years, 11 months ago

Answer : Variant and Array https://docs.snowflake.com/en/user-guide/semistructured-intro.html#loading-semi-structured-data

upvoted 2 times

Question #247 *Topic 1*

In an auto-scaling multi-cluster virtual warehouse with the setting SCALING_POLICY = ECONOMY enabled, when is another cluster started?

A. When the system has enough load for 2 minutes

B. When the system has enough load for 6 minutes

C. When the system has enough load for 8 minutes

D. When the system has enough load for 10 minutes

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B. When the system has enough load for 6 minutes

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. When the system has enough load for 6 minutes

upvoted 1 times

**singhks** 3 years, 2 months ago

B is the correct answer

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**v1tjkee** 3 years, 4 months ago

What's wrong with these questions? the answer is 6 minutes

upvoted 1 times

**srj_mehta** 3 years, 4 months ago

Only if the system estimates there's enough query load to keep the cluster busy for at least 6 minutes. B.

upvoted 1 times

**OscarWest2023** 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/warehouses-multicluster#setting-the-scaling-policy-for-a-multi-cluster-warehouse

upvoted 1 times

**fahfouhi94** 3 years, 9 months ago

https://docs.snowflake.com/en/user-guide/warehouses-multicluster.html warehouse start Only if the system estimates there's enough query load to keep the cluster busy for at least 6 minutes.

upvoted 1 times

**nintendogamer64** 3 years, 10 months ago

B: Very straightforward in the documentation

upvoted 3 times

**AlexbDku** 3 years, 10 months ago

6 minutes https://docs.snowflake.com/en/user-guide/warehouses-multicluster.html

upvoted 4 times

**miadiesel** 3 years, 10 months ago

Only if the system estimates there's enough query load to keep the cluster busy for at least 6 minutes https://docs.snowflake.com/en/user-guide/warehouses-multicluster.html

upvoted 4 times

**harshagc180** 3 years, 11 months ago

Answer B https://docs.snowflake.com/en/user-guide/warehouses-multicluster.html

upvoted 3 times

Question #248 *Topic 1*

What is the following SQL command used for?

Select * from table(validate(t1, job_id => '_last'));

A. To validate external table files in table t1 across all sessions

B. To validate task SQL statements against table t1 in the last 14 days

C. To validate a file for errors before it gets executed using a COPY command

D. To return errors from the last executed COPY command into table t1 in the current session

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

D is the correct answer

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. To return errors from the last executed COPY command into table t1 in the current session

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

correct

upvoted 1 times

**SV1122** 3 years, 10 months ago

Validates the files loaded in a past execution of the COPY INTO command and returns all the errors encountered during the load, rather than just the first error. https://docs.snowflake.com/en/sql-reference/functions/validate.html

upvoted 3 times

**halol** 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/functions/validate.html

upvoted 1 times

Question #249 *Topic 1*

A sales table FCT_SALES has 100 million records.

The following query was executed:

SELECT COUNT (1) FROM FCT_SALES;

How did Snowflake fulfill this query?

A. Query against the result set cache

B. Query against a virtual warehouse cache

C. Query against the most-recently created micro-partition

D. Query against the metadata cache

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**kolossal** Most Recent 1 year, 12 months ago

The COUNT(1) function is used to count all rows in the specified table. In Snowflake, using COUNT(1) is equivalent to using COUNT(*), as both provide a count of all rows without filtering based on any specific column. The number 1 is simply a placeholder; it does not affect the result. Snowflake optimizes this operation by using metadata statistics, allowing for quick retrieval of row counts without scanning every row.

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

D.Query against the metadata cache

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D is correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**EmiB** 3 years, 7 months ago

D. Query against the metadata cache >> METADATA-BASED RESULT

upvoted 2 times

**SV1122** 3 years, 10 months ago

D. Query against the metadata cache

upvoted 1 times

Question #250 *Topic 1*

What happens when a virtual warehouse is resized?

A. When increasing the size of an active warehouse the compute resource for all running and queued queries on the warehouse are affected.

B. When reducing the size of a warehouse the compute resources are removed only when they are no longer being used to execute any current statements.

C. The warehouse will be suspended while the new compute resource is provisioned and will resume automatically once provisioning is complete.

D. Users who are trying to use the warehouse will receive an error message until the resizing is complete.

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**halol** Highly Voted 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/warehouses-tasks.html#:~:text=Resizing%20a%20running%20warehouse%20adds,each%20cluster%20in%20the%20warehouse.

upvoted 13 times

**Kate0204** Highly Voted 3 years, 11 months ago

B. When reducing the size of a warehouse the compute resources are removed only when they are no longer being used to execute any current statements.

upvoted 7 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B. When reducing the size of a warehouse the compute resources are removed only when they are no longer being used to execute any current statements.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. When reducing the size of a warehouse the compute resources are removed only when they are no longer being used to execute any current statements.

upvoted 1 times

**Afzy** 3 years ago

I think answer is D, Warehouse Restart: Depending on the resizing operation, Snowflake may need to restart the virtual warehouse. During this restart, the warehouse is temporarily unavailable for query processing.

upvoted 1 times

**singhks** 3 years, 2 months ago

B is correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 2 times

**Needium** 3 years, 6 months ago

B is the answer

upvoted 1 times

Question #251 *Topic 1*

What tasks can be completed using the COPY command? (Choose two.)

A. Columns can be aggregated.

B. Columns can be joined with an existing table.

C. Columns can be reordered.

D. Columns can be omitted.

E. Data can be loaded without the need to spin up a virtual warehouse.

**Suggested Answer:** *C, D*

*Community vote distribution*

CD (100%)

**Mallikharjuna452** 2 years, 3 months ago

columns can be reordered and omitted

upvoted 1 times

**KarBiswa** Highly Voted 4 months ago

Selected Answer: C

Query Processing

upvoted 10 times

**_yyukta** 2 years, 8 months ago

C and D are correct

upvoted 1 times

**dmitriypo** 4 months ago

Selected Answer: A

This is a tricky question. Authentication, Access Control, and Security can be configured. And these are the parts of the Cloud Services layer. On the other hand Query processing layer consists of Virtual Data Warehouses which also can be configured (size and auto-scaling options). But I would most probably go for A.

upvoted 4 times

**singhks** 3 years, 2 months ago

C and D

upvoted 1 times

**chatrisrinu** 4 months ago

Selected Answer: B

I think answer should be B (Cloud Services) layer can be configured. Ex: Cloud Services layer to use a different network topology. For example, you can configure Snowflake to use a virtual private cloud (VPC) to isolate your data from other users on the Snowflake platform.

upvoted 2 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**JG1984** Most Recent 4 months ago

Snowflake manages all aspects of how the data is stored and organized, and the user cannot modify them.

upvoted 2 times

**c0d3g** 3 years, 5 months ago

Columns can be reordered. Columns can be omitted.

upvoted 1 times

**mayank** 4 months ago

Selected Answer: C

C- query processing

upvoted 2 times

**KarBiswa** 3 years, 8 months ago

Right answer

upvoted 1 times

**girgir** 3 years, 8 months ago

yes CD, Note that COPY transformations do not support the FLATTEN function, or JOIN or GROUP BY (aggregate) syntax.
https://docs.snowflake.com/en/user-guide/data-load-transform

upvoted 4 times

**fahfouhi94** 3 years, 9 months ago

https://docs.snowflake.com/en/user-guide/data-load-transform.html

upvoted 2 times

Question #252 *Topic 1*

Which Snowflake layer can be configured?

A. Database Storage

B. Cloud Services

C. Query Processing

D. Application Services

**Suggested Answer:** *C*

*Community vote distribution*

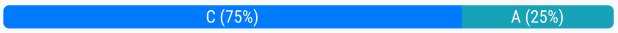


**KarBiswa** Highly Voted 3 years, 8 months ago

Query Processing

upvoted 10 times

**dmitriypo** Highly Voted 3 years, 5 months ago

This is a tricky question. Authentication, Access Control, and Security can be configured. And these are the parts of the Cloud Services layer. On the other hand Query processing layer consists of Virtual Data Warehouses which also can be configured (size and auto-scaling options). But I would most probably go for A.Snowflake manages all aspects of how the data is stored and organized, and the user cannot modify them.

upvoted 5 times

**JG1984** 3 years, 1 month ago

Snowflake manages all aspects of how the data is stored and organized, and the user cannot modify them.

upvoted 2 times

**Goransoll** Most Recent 7 months, 4 weeks ago

Query Processing - users configure warehouse size, scaling, suspend/resume, etc. No: A - Database storage - Snowflake manages all storage automatically

upvoted 1 times

**bor4un** 1 year, 9 months ago

Seems correct - warehouses

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

Query Processing

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

Query processing layer consists of Virtual Data Warehouses which also can be configured (size and auto-scaling options).

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C. Query processing

upvoted 1 times

**chatrisrinu** 3 years ago

I think answer should be B (Cloud Services) layer can be configured. Ex: Cloud Services layer to use a different network topology. For example, you can configure Snowflake to use a virtual private cloud (VPC) to isolate your data from other users on the Snowflake platform.

upvoted 2 times

**dmitriypo** 3 years, 5 months ago

Selected Answer: A

This is a tricky question. Authentication, Access Control, and Security can be configured. And these are the parts of the Cloud Services layer. On the other hand Query processing layer consists of Virtual Data Warehouses which also can be configured (size and auto-scaling options). But I would most probably go for A.

upvoted 4 times

**singhks** 3 years, 2 months ago

C is correct

upvoted 1 times

**andretala** 3 years, 3 months ago

A- Database Storage

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

correct

upvoted 1 times

**weuz** 3 years, 7 months ago

the C option is the correct one

upvoted 1 times

**mayank** 3 years, 8 months ago

C- query processing

upvoted 2 times

**Oldmandonut** 3 years, 8 months ago

A, B, C three layer is right

upvoted 1 times

**markpumc** 3 years, 9 months ago

application Service is not a snowflake layer, should be C

upvoted 2 times

**lawgladvin** 3 years, 8 months ago

Answer is C

upvoted 1 times

**Andywu213** 3 years, 10 months ago

C, query processing

upvoted 2 times

**Naveen90** 3 years, 10 months ago

I think the Answer is C - Query processing. Warehouse can be configured, not sure about Cloud services or Storage(Snowflake manages this)

upvoted 2 times

Question #253 *Topic 1*

Query compilation occurs in which architecture layer of the Snowflake Cloud Data Platform?

A. Compute layer

B. Storage layer

C. Cloud infrastructure layer

D. Cloud services layer

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**halol** Highly Voted 3 years, 10 months ago

D. Cloud services layer https://community.snowflake.com/s/article/Snowflake-Additional-Billing-Cloud-Services-Layer-Impact#:~:text=The%20services%20layer%20for%20Snowflake,cache%20and%20coordinates%20all%20transactions.

upvoted 10 times

**isolated_sniper** Highly Voted 3 years, 6 months ago

Ans - D. Query execution is different from Query processing. Query execution is performed in the processing layer. While The services layer for Snowflake authenticates user sessions, provides management, enforces security functions, performs query compilation and optimization, results cache and coordinates all transactions

upvoted 6 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Cloud services layer

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

D. Cloud services layer The Cloud Services layer in Snowflake is responsible for activities such as query optimization, compilation, and overall management of the system.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D is correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

correct

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

D. Cloud services layer https://community.snowflake.com/s/article/Snowflake-Additional-Billing-Cloud-Services-Layer-Impact#:~:text=The%20services%20layer%20for%20Snowflake,cache%20and%20coordinates%20all%20transactions.

upvoted 2 times

**fahfouhi94** 3 years, 9 months ago

cloud services layer

upvoted 1 times

**BigDataBB** 3 years, 10 months ago

https://docs.snowflake.com/en/sql-reference/sql/explain.html Although EXPLAIN does not consume any compute credits, the compilation of the query does consume Cloud Service credits, just as other metadata operations do.

upvoted 2 times

**BigDataBB** 3 years, 10 months ago

https://docs.snowflake.com/en/sql-reference/sql/explain.html Although EXPLAIN does not consume any compute credits, the compilation of the query does consume Cloud Service credits, just as other metadata operations do.https://docs.snowflake.com/en/user-guide/cost-understanding-compute.html#cloud-service-credit-usage

upvoted 2 times

**BigDataBB** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/cost-understanding-compute.html#cloud-service-credit-usage

upvoted 1 times

**kokosz** 3 years, 10 months ago

services

upvoted 1 times

**AlexbDku** 3 years, 10 months ago

Answer is D

upvoted 3 times

**vishgan** 3 years, 10 months ago

D is right

upvoted 2 times

Question #254 *Topic 1*

If a size Small virtual warehouse is made up of two servers, how many servers make up a Large warehouse?

A. 4

B. 8

C. 16

D. 32

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**LinoT** Highly Voted 3 years, 8 months ago

Ok, after 250+ questions so far out of 600+ I have to say I am very disappointed in the quality of the answers. 70% of the answers provided by Examtopics are incorrect in this exam prep. I expect to have 2% or so wrong answers but 70% is unacceptable is causing more confusion that helping at this point. For a paying customer, I would say that is absolutely unacceptable.95% trust to MOST VOTEDNo one should pay this. Just scroll through the pages on free.rightly said- least not expected wrong answers for these damn direct questionsHaving most of the answer wrong helps people to study the comments section, go through the weblinks. Having 100% correct answers will only make people dumb. They will just memorize the question and answers. Its a risk. Having wrong answers helped me to test my knowledge.

upvoted 41 times

**LinoT** Highly Voted 3 years, 8 months ago

Ok, after 250+ questions so far out of 600+ I have to say I am very disappointed in the quality of the answers. 70% of the answers provided by Examtopics are incorrect in this exam prep. I expect to have 2% or so wrong answers but 70% is unacceptable is causing more confusion that helping at this point. For a paying customer, I would say that is absolutely unacceptable.

upvoted 41 times

**EmiB** 3 years, 7 months ago

95% trust to MOST VOTED

upvoted 10 times

**AkexS** 3 years, 3 months ago

No one should pay this. Just scroll through the pages on free.

upvoted 1 times

**Matsalearnings** 3 years, 1 month ago

rightly said- least not expected wrong answers for these damn direct questions

upvoted 2 times

**Rajivnb** 3 years ago

Having most of the answer wrong helps people to study the comments section, go through the weblinks. Having 100% correct answers will only make people dumb. They will just memorize the question and answers. Its a risk. Having wrong answers helped me to test my knowledge.

upvoted 4 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

8 is correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B is correct

upvoted 1 times

**bbrandhh** 3 years, 3 months ago

this question was on my exam july 2023

upvoted 2 times

**examtopics_strata** 3 years, 3 months ago

Correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**rneti** 3 years, 7 months ago

Per the Snowflake guide (https://docs.snowflake.com/en/user-guide/warehouses-overview.html) the answer should be (B) 8 Credits. The answer provided by the moderator is wrong. Couldn't complain as this is a free service (:

upvoted 4 times

**KarBiswa** 3 years, 8 months ago

8 is the number

upvoted 1 times

**ntechkuma** 3 years, 8 months ago

Its B. Moderator should update correct answer. Most if the questions have wrong answers

upvoted 1 times

**fahfouhi94** 3 years, 9 months ago

https://docs.snowflake.com/en/user-guide/warehouses-overview.html

upvoted 2 times

**AlexbDku** 3 years, 10 months ago

Answer is B https://docs.snowflake.com/en/user-guide/warehouses-overview.html#warehouse-size

upvoted 4 times

**miadiesel** 3 years, 10 months ago

Small 2; Medium 4; Large 8; X-Large 16; https://docs.snowflake.com/en/user-guide/warehouses-overview.html#warehouse-size

upvoted 3 times

**SnowProCertDec22** 3 years, 10 months ago

answer b

upvoted 1 times

**Kate0204** 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/warehouses-overview.html#warehouse-size

upvoted 2 times

Question #255 *Topic 1*

A clustering key was defined on a table, but it is no longer needed.

How can the key be removed?

A. ALTER TABLE [TABLE NAME] PURGE CLUSTERING KEY

B. ALTER TABLE [TABLE NAME] DELETE CLUSTERING KEY

C. ALTER TABLE [TABLE NAME] DROP CLUSTERING KEY

D. ALTER TABLE [TABLE NAME] REMOVE CLUSTERING KEY

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**AS314** Highly Voted 3 years, 10 months ago

ALTER TABLE DROP CLUSTERING KEY

upvoted 10 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

C is correct answer

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C is correct

upvoted 1 times

**Amitsnowflake** 3 years, 1 month ago

C https://docs.snowflake.com/en/user-guide/tables-clustering-keys

upvoted 1 times

**dmitriypo** 3 years, 5 months ago

C https://docs.snowflake.com/en/user-guide/tables-clustering-keys#defining-clustered-tables

upvoted 1 times

**GiselaS** 3 years, 7 months ago

Dropping the Clustering Keys for a Table At any time, you can drop the clustering key for a table using ALTER TABLE: ALTER TABLE DROP CLUSTERING KEY

upvoted 2 times

**weuz** 3 years, 7 months ago

C is the correct one

upvoted 2 times

**KarBiswa** 3 years, 8 months ago

Alter table drop clustering key

upvoted 2 times

**fahfouhi94** 3 years, 9 months ago

https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html

upvoted 1 times

**halol** 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/sql/alter-table.html C

upvoted 3 times

**paknadeem** 3 years, 11 months ago

I Think the correct answer is C

upvoted 2 times

Question #256 *Topic 1*

What is a core benefit of clustering?

A. To guarantee uniquely identifiable records in the database

B. To increase scan efficiency in queries by improving pruning

C. To improve performance by creating a separate file for point lookups

D. To provide data redundancy by duplicating micro-partitions

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B is correct answer

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. To increase scan efficiency in queries by improving pruning

upvoted 1 times

**Marge23** 3 years, 1 month ago

correct

upvoted 1 times

**fahfouhi94** 3 years, 9 months ago

correct answer

upvoted 3 times

**AlexbDku** 3 years, 10 months ago

B is correct https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html#:~:text=Using%20a%20clustering%20key%20to,in%20tables%20with%20no%20clustering.

upvoted 4 times

Question #257 *Topic 1*

Which statement is true about Multi-Factor Authentication (MFA) in Snowflake?

A. MFA can be enforced or applied for a given role.

B. Snowflake users are automatically enrolled in MFA.

C. Users enroll in MFA by submitting a request to Snowflake Support.

D. MFA is an integrated Snowflake feature.

**Suggested Answer:** *D*

*Community vote distribution*

D (98%) 2

**AlexbDku** Highly Voted 3 years, 4 months ago

D is right https://docs.snowflake.com/en/user-guide/security-mfa.html#:~:text=Snowflake%20supports%20multi%2Dfactor%20authentication,is%20managed%20completely%20by%20Snowflake.MFA is enabled on a per-user basis; however, at this time, users are not automatically enrolled in MFA. To use MFA, users must enroll themselves.Yes, though MFA is still an integrated feature for Snowflake, just not activated for every user automatically.

upvoted 13 times

**AlexbDku** Highly Voted 3 years, 4 months ago

Selected Answer: D

D is right https://docs.snowflake.com/en/user-guide/security-mfa.html#:~:text=Snowflake%20supports%20multi%2Dfactor%20authentication,is%20managed%20completely%20by%20Snowflake.

upvoted 13 times

**BigDataBB** 3 years, 3 months ago

MFA is enabled on a per-user basis; however, at this time, users are not automatically enrolled in MFA. To use MFA, users must enroll themselves.Yes, though MFA is still an integrated feature for Snowflake, just not activated for every user automatically.

upvoted 1 times

**BigDataBB** 3 years, 3 months ago

MFA is enabled on a per-user basis; however, at this time, users are not automatically enrolled in MFA. To use MFA, users must enroll themselves.

upvoted 1 times

**SexyPotato** 3 years, 2 months ago

Yes, though MFA is still an integrated feature for Snowflake, just not activated for every user automatically.

upvoted 1 times

**shaikhyasar** Most Recent 1 year, 8 months ago

A is wrong because MFA is applied to user level not row level

upvoted 2 times

**_yyukta** 2 years, 2 months ago

A. MFA can be enforced or applied for a given role.

upvoted 1 times

**lol632** 2 years, 5 months ago

D is correct

upvoted 1 times

**GAUURAAV** 2 years, 5 months ago

D is correct

upvoted 1 times

**matias_salina** 2 years, 7 months ago

https://docs.snowflake.com/en/user-guide/security-mfa

upvoted 1 times

**Amitsnowflake** 2 years, 7 months ago

https://docs.snowflake.com/en/user-guide/security-mfa D is correct

upvoted 1 times

**SnowYeti** 2 years, 7 months ago

D is correct

upvoted 1 times

**yourteacher** 2 years, 8 months ago

D is correct

upvoted 2 times

**examtopics_strata** 2 years, 9 months ago

Correct

upvoted 1 times

**Leinho** 2 years, 10 months ago

https://docs.snowflake.com/en/user-guide/security-mfa.html

upvoted 1 times

**srj_mehta** 2 years, 10 months ago

MFA support is provided as an integrated Snowflake feature, powered by the Duo Security service, which is managed completely by Snowflake. MFA is enabled on a per-user basis; however, at this time, users are not automatically enrolled in MFA. To use MFA, users must enroll themselves.

https://docs.snowflake.com/en/user-guide/security-mfa

upvoted 2 times

**Needium** 2 years, 12 months ago

D s the right answer

upvoted 1 times

**Juciifer** 3 years ago

D is correct

upvoted 1 times

**Raju039** 3 years, 1 month ago

D is right

upvoted 2 times

**EmiB** 3 years, 1 month ago

D. MFA is an integrated Snowflake feature.

upvoted 3 times

**KarBiswa** 3 years, 1 month ago

best of the options

upvoted 2 times

Question #258 *Topic 1*

What data type should be used to store JSON data natively in Snowflake?

A. JSON

B. String

C. Object

D. VARIANT

**Suggested Answer:** *D*

*Community vote distribution*

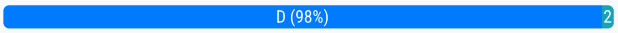


**oldcrowtheory** Highly Voted 3 years, 7 months ago

Variant data type

upvoted 11 times

**halol** Highly Voted 3 years, 11 months ago

D is correct

upvoted 5 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Variant data type

upvoted 1 times

**Pre_Nan_123** 2 years, 8 months ago

NO JSON TYPE. Variant is right answer.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D is correct

upvoted 1 times

**AlphaCentauriBear** 2 years, 8 months ago

There is no "JSON" data type. https://docs.snowflake.com/en/user-guide/script-data-load-transform-json When uploading JSON data into a table, you have these options: Store JSON objects natively in a VARIANT type column

upvoted 1 times

**e_l_i_a_s** 2 years, 9 months ago

D is correct

upvoted 1 times

**azlovemovie** 2 years, 10 months ago

Correct answer is D. Variant

upvoted 1 times

**Bhargava12** 2 years, 11 months ago

Variant data type

upvoted 1 times

**GAUURAAV** 2 years, 11 months ago

D - Variant

upvoted 1 times

**Avi0407** 2 years, 12 months ago

D And that is basic, I feel.... Why are you mentioning wrong answers most of the time, Team Examtopics ???

upvoted 1 times

**Afzy** 3 years ago

Answer is D

upvoted 1 times

**Gaurangvp44** 3 years ago

DDDDDDDDDDDDDD https://www.bing.com/search?pglt=41&q=What+data+type+should+be+used+to+store+JSON+data+natively+in+Snowflake%3F&cvid=680d812119164b8788a54e9fa50557a8&gs_lcrp=EgZjaHJvbWUyBggAEEUYOdIBBzYzNmowajGoAgCwAgA&FORM=ANNTA1&PC=DCTS

upvoted 1 times

**vrfidman** 3 years ago

The only possible option.

upvoted 1 times

**Heetec** 3 years, 1 month ago

D - variant

upvoted 1 times

**SnowYeti** 3 years, 1 month ago

CD are correct!

upvoted 1 times

**timcooks** 3 years, 1 month ago

Varient

upvoted 1 times

Question #259 *Topic 1*

What should be considered when deciding to use a Secure View? (Choose two.)

A. No details of the query execution plan will be available in the query profiler.

B. Once created there is no way to determine if a view is secure or not.

C. Secure views do not take advantage of the same internal optimizations as standard views.

D. It is not possible to create secure materialized views.

E. The view definition of a secure view is still visible to users by way of the information schema.

**Suggested Answer:** *A, C*

*Community vote distribution*




**AS314** Highly Voted 3 years, 10 months ago

https://alexandersks.medium.com/views-in-snowflake-part-3-secure-view-8bbd8832b3af

upvoted 13 times

**halol** Highly Voted 3 years, 11 months ago

A & C i think

upvoted 8 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A and C are correct

upvoted 1 times

**Vijay2605** 2 years, 8 months ago

AC is the right ans

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A and C are correct

upvoted 1 times

**GAUURAAV** 2 years, 11 months ago

A and C

upvoted 1 times

**Avi0407** 2 years, 12 months ago

https://alexandersks.medium.com/views-in-snowflake-part-3-secure-view-8bbd8832b3af

upvoted 2 times

**Heetec** 3 years, 1 month ago

A and C

upvoted 1 times

**Thaiya** 3 years, 1 month ago

A and C

upvoted 1 times

**Amitsnowflake** 3 years, 2 months ago

A and C

upvoted 1 times

**HighBMI** 3 years, 5 months ago

AC are the correct answers

upvoted 1 times

**Needium** 3 years, 6 months ago

AC is correct

upvoted 1 times

**aroma** 3 years, 6 months ago

AC is correct

upvoted 2 times

**Raju039** 3 years, 7 months ago

The internals of a secure view are not exposed in Query Profile (in the web interface).Some of the internal optimizations for views require access to the underlying data in the base tables for the view. This access might allow data that is hidden from users of the view to be exposed through user code, such as user-defined functions, or other programmatic methods. Secure views do not utilize these optimizations, ensuring that users have no access to the underlying data.

upvoted 4 times

**Raju039** 3 years, 7 months ago

**Selected Answer: AC**

The internals of a secure view are not exposed in Query Profile (in the web interface).

upvoted 4 times

**Raju039** 3 years, 7 months ago

Some of the internal optimizations for views require access to the underlying data in the base tables for the view. This access might allow data that is hidden from users of the view to be exposed through user code, such as user-defined functions, or other programmatic methods. Secure views do not utilize these optimizations, ensuring that users have no access to the underlying data.

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

correct answer

upvoted 2 times

**mayank** 3 years, 8 months ago

a and c.

upvoted 2 times

**sailoo** 3 years, 9 months ago

A,C for sure

upvoted 1 times

Question #260 *Topic 1*

The information schema provides storage information for which of the following objects? (Choose two.)

A. Users

B. Databases

C. Internal stages

D. Resource monitors

E. Pipes

**Suggested Answer:** *B, C*

*Community vote distribution*




**Kate0204** Highly Voted 3 years, 11 months ago

Each database created in your account automatically includes a built-in, read-only schema named INFORMATION_SCHEMA. The schema contains the following objects:

Views for all the objects contained in the database, as well as views for account-level objects (i.e. non-database objects such as roles, warehouses, and databases)

Table functions for historical and usage data across your account.

upvoted 7 times

**Kate0204** Highly Voted 3 years, 11 months ago

Each database created in your account automatically includes a built-in, read-only schema named INFORMATION_SCHEMA. The schema contains the following objects: Views for all the objects contained in the database, as well as views for account-level objects (i.e. non-database objects such as roles, warehouses, and databases) Table functions for historical and usage data across your account.Correctly, every DATABASE has his own INFORMATION_SCHEMA so A INFORMATION_SCHEMA has information about object of a SINGLe database, in this case: Internal stages Pipes So B, C

upvoted 7 times

**BigDataBB** 3 years, 9 months ago

Correctly, every DATABASE has his own INFORMATION_SCHEMA so A INFORMATION_SCHEMA has information about object of a SINGLe database, in this case: Internal stages Pipes So B, C

upvoted 2 times

**SV1122** 3 years, 10 months ago

Selected Answer: BE

While the reasoning for the Database is clear. I don't understand how to choose between internal stages vs pipes. The Info Schema has views offering info on both stages and pipes.

upvoted 6 times

**SV1122** 3 years, 10 months ago

While the reasoning for the Database is clear. I don't understand how to choose between internal stages vs pipes. The Info Schema has views offering info on both stages and pipes.The question is about storage information, there is no storage info about pipes. So the correct answer is only BC.BCE are correct

upvoted 6 times

**aemilka** Most Recent 2 years, 1 month ago

The question is about storage information, there is no storage info about pipes. So the correct answer is only BC.

upvoted 2 times

**EmiB** 3 years, 7 months ago

BCE are correct

upvoted 3 times

**Mallikharjuna452** 2 years, 3 months ago

Databases and internal stages

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B and C are correct

upvoted 1 times

**Afzy** 3 years ago

Databases,Schemas,Tables,Columns,Views: Stages: File Formats: Warehouses: Users and Roles: Privileges: Sequences: Resource Monitors:

upvoted 1 times

**Afzy** 3 years ago

Right answer is B,C, and D but they are asking for only two

upvoted 1 times

**Rajivnb** 3 years ago

Though Information_Schema provides the Pipes details, the question is about "Storage" Information. So only DB and Internal Stages will have Storage related information. If "Storage " is not used in the question, then Pipe is the 3rd answer.

upvoted 3 times

**nonamae** 3 years ago

BC - based on the documentation, information schema displays view for database and account level objects as well as views for non-standard objects like Stages and File Formats. https://docs.snowflake.com/en/sql-reference/info-schema

upvoted 1 times

**Marge23** 3 years, 1 month ago

definitions of databases and pipes. no internal stages.

upvoted 1 times

**singhks** 3 years, 2 months ago

DATABASE_STORAGE_USAGE_HISTORY STAGE_STORAGE_USAGE_HISTORY TABLE_STORAGE_METRICS

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**SteveCrown** 3 years, 5 months ago

The answer is B & E https://docs.snowflake.com/en/sql-reference/info-schema information_schema contains views that hold metadata for - Databases, Stages(not Internal stages even though internal stage are examples of stages) and pipes. Therefore Databases and Pipes are the correct options.

upvoted 4 times

**SteveCrown** 3 years, 5 months ago

The answer is B & E https://docs.snowflake.com/en/sql-reference/info-schema information_schema contains views that hold metadata for - Databases, Stages(not Internal stages even though internal stage are examples of stages) and pipes. Therefore Databases and Pipes are the correct options.correct ! i can see "DEMO_DB / INFORMATION_SCHEMA / PIPES" in information schema

upvoted 4 times

**NEMMELA** 3 years, 1 month ago

correct ! i can see "DEMO_DB / INFORMATION_SCHEMA / PIPES" in information schema

upvoted 1 times

**dmitriypo** 3 years, 5 months ago

BCE should be correct, but I don't see any storage information in those views.

upvoted 1 times

**isolated_sniper** 3 years, 6 months ago

My opinion - BC E - Information schema gives PIPE usage information, not storage information.

upvoted 2 times

**Kvk117** 3 years, 7 months ago

B,C,E is the correct answer. Verified it

upvoted 2 times

**Raju039** 3 years, 7 months ago

B, C, E are correct answers.

upvoted 1 times

**EmiB** 3 years, 8 months ago

Old question - 3 are correct: B. Databases - DATABASES C. Internal stages - STAGES E. Pipes - PIPESDoes pipe store any data? The question is asking about "storage information"

upvoted 6 times

**EmiB** 3 years, 8 months ago

**Selected Answer: BC**

Old question - 3 are correct: B. Databases - DATABASES C. Internal stages - STAGES E. Pipes - PIPES

upvoted 6 times

**0e504b5** 2 years, 9 months ago

Does pipe store any data? The question is asking about "storage information"

upvoted 1 times

Question #261 *Topic 1*

What is a responsibility of Snowflake's virtual warehouses?

A. Infrastructure management

B. Metadata management

C. Query execution

D. Query parsing and optimization

E. Management of the storage layer

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**_yyukta** Most Recent 2 years, 2 months ago

C is correct

upvoted 1 times

**Marge23** 2 years, 7 months ago

correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**EmiB** 3 years, 1 month ago

C. Query execution

upvoted 2 times

**vmsaran** 3 years, 2 months ago

Ans is C

upvoted 2 times

Question #262 *Topic 1*

Which data type is supported by Snowflake data classification?

A. Binary

B. Float

C. Geography

D. Variant

**Suggested Answer:** *B*

*Community vote distribution*




**halol** Highly Voted 3 years, 11 months ago

https://docs.snowflake.com/en/user-guide/governance-classify-concepts.html#what-is-classification You can classify table and view columns for all supported data types except for the following data types: GEOGRAPHY BINARY VARIANT

upvoted 13 times

**Raju039** Highly Voted 4 months ago

Selected Answer: C

LZO is for Parquet file and GZIP is for JSON file.

upvoted 9 times

**niho1993** 1 year, 1 month ago

Snowflake Data Classification supports structured and semi-structured data types, allowing automated sensitivity detection (e.g., PII data). Among the given options, only VARIANT is supported for classification

upvoted 1 times

**aroma** 4 months ago

TYPE = PARQUET COMPRESSION = AUTO | LZO | SNAPPY | NONE SNAPPY Files are compressed using the Snappy algorithm by default. You can optionally specify this value.

upvoted 4 times

**I_Rehan_1** 1 year, 7 months ago

Snowflake Data Classification supports structured and semi-structured data types, allowing automated sensitivity detection (e.g., PII data). Among the given options, only VARIANT is supported for classification

upvoted 2 times

**EmiB** 4 months ago

Selected Answer: C

LZO and SNAPPY

upvoted 3 times

**Kalia_ipsita** 1 year, 11 months ago

You can classify table and view columns for all supported data types except for the following data types: ARRAY, BINARY,GEOGRAPHY,OBJECT,VARIANT ( Note that you can classify a column with the VARIANT data type when the column data type can be cast to a NUMBER or STRING data type. Snowflake does not classify the column if the column contains JSON, XML, or other semi-structured data.), VECTOR If a table contains columns that are not of a supported data type or the column contains all NULL values, the classification process ignores the columns and does not include them in the output. Important If your data represents NULL values with a value other than NULL, the accuracy of the classification results may be impacted. https://docs.snowflake.com/user-guide/classify-intro#supported-objects-and-data-types

upvoted 2 times

**SexyPotato** Most Recent 4 months ago

None if the above. by default it should be Snappy. "When unloading data, files are compressed using the Snappy algorithm by default. You can optionally specify this value." https://docs.snowflake.com/en/sql-reference/sql/create-file-format#type-parquet:~:text=When%20unloading%20data%2C%20files%20are%20compressed%20using%20the%20Snappy%20algorithm%20by%20default.%20You%20can%20optionally%20specify%20this%20value.

upvoted 3 times

**kolossal** 1 year, 12 months ago

except for these data type ARRAY, BINARY, GEOGRAPHY, OBJECT, VARIANT, VECTOR.

upvoted 2 times

**KarBiswa** 4 months ago

Selected Answer: C

Parquet compression

upvoted 2 times

**Rameez1** 2 years, 4 months ago

https://docs.snowflake.com/en/user-guide/governance-classify-concepts.html#what-is-classification

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B.Float

upvoted 1 times

**Wzghaer** 2 years, 11 months ago

B is the right answer https://docs.snowflake.com/en/user-guide/governance-classify-concepts

upvoted 1 times

**rakshithaK** 3 years, 3 months ago

https://docs.snowflake.com/en/user-guide/governance-classify-concepts

upvoted 1 times

**examtopics_strata** 3 years, 3 months ago

Float is correct

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

GEOGRAPHY BINARY VARIANT

upvoted 1 times

**AlexbDku** 3 years, 10 months ago

Answer is float https://docs.snowflake.com/en/user-guide/governance-classify-concepts.html#what-is-classification

upvoted 3 times

**harshagc180** 3 years, 11 months ago

As per the below link, answer should be Float. Please confirm. https://docs.snowflake.com/en/user-guide/governance-classify-concepts.html

upvoted 4 times

Question #263 *Topic 1*

When unloading data to an external stage, which compression format can be used for Parquet files with the COPY INTO command?

A. BROTLI

B. GZIP

C. LZO

D. ZSTD

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**Raju039** Highly Voted 3 years, 7 months ago

LZO is for Parquet file and GZIP is for JSON file.

upvoted 9 times

**PrashantGupta1616** Most Recent 2 years, 5 months ago

LZO is for Parquet file and GZIP is for JSON file.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

C. LZO

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**dmitriypo** 3 years, 5 months ago

LZO https://docs.snowflake.com/en/sql-reference/sql/copy-into-location#syntax

upvoted 1 times

**aroma** 3 years, 6 months ago

TYPE = PARQUET COMPRESSION = AUTO | LZO | SNAPPY | NONE SNAPPY Files are compressed using the Snappy algorithm by default. You can optionally specify this value.

upvoted 4 times

**KarBiswa** 3 years, 8 months ago

Parquet compression

upvoted 2 times

**EmiB** 3 years, 8 months ago

LZO and SNAPPY

upvoted 3 times

**SexyPotato** 3 years, 8 months ago

None if the above. by default it should be Snappy. "When unloading data, files are compressed using the Snappy algorithm by default. You can optionally specify this value." https://docs.snowflake.com/en/sql-reference/sql/create-file-format#type-parquet:~:text=When%20unloading%20data%2C%20files%20are%20compressed%20using%20the%20Snappy%20algorithm%20by%20default.%20You%20can%20optionally%20specify%20this%20value.

upvoted 3 times

**thejeby** 3 years, 9 months ago

C - LZO: https://docs.snowflake.com/en/sql-reference/sql/copy-into-location.html#type-parquet Options Default: Auto = Snappy LZO Snappy None Other options are:

upvoted 2 times

**BigDataBB** 3 years, 9 months ago

if You try You will disvover that "By default, all unloaded data files are compressed using gzip" is false for parquet file type COPY INTO @SG_PROVA FROM GAME_LOGS FILE_FORMAT = ( TYPE = PARQUET); -- sg_prova/data_0_0_0.snappy.parquet COPY INTO @SG_PROVA/data2 FROM GAME_LOGS FILE_FORMAT = ( TYPE = PARQUET COMPRESSION = LZO); -- sg_prova/data2_0_0_0.lzo.parquet COPY INTO @SG_PROVA/data3 FROM GAME_LOGS FILE_FORMAT = ( TYPE = PARQUET COMPRESSION = GZIP); -- Invalid compression format GZIP

upvoted 2 times

**SV1122** 3 years, 10 months ago

LZO https://docs.snowflake.com/en/sql-reference/sql/copy-into-location.html#format-type-options-formattypeoptions

upvoted 1 times

**SD1SD** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/intro-summary-unloading.html Compression of Output Data Files The following table describes how Snowflake handles compression for the output files generated by Snowflake when unloading data: Location of Files Supported Notes Internal or external location gzip By default, all unloaded data files are compressed using gzip, unless compression is explicitly disabled or one of the other supported compression methods is explicitly specified. bzip2 Brotli Zstandard

upvoted 1 times

**AlexbDku** 3 years, 10 months ago

C - LZO

upvoted 2 times

**halol** 3 years, 11 months ago

When unloading data, files are compressed using the Snappy algorithm by default. If unloading data to LZO-compressed files, specify this value. https://docs.snowflake.com/en/sql-reference/sql/create-file-format.html#type-parquet

upvoted 2 times

Question #264 *Topic 1*

Which SQL command can be used to verify the privileges that are granted to a role?

A. SHOW GRANTS ON ROLE [Role Name]

B. SHOW ROLES [Role Name]

C. SHOW GRANTS TO ROLE [Role Name]

D. SHOW GRANTS FOR ROLE [Role Name]

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**miadiesel** Highly Voted 2 years, 10 months ago

Answer shown does not even exist. Show grants to role: List all privileges granted https://docs.snowflake.com/en/sql-reference/sql/show-grants.html#show-grants

upvoted 6 times

**romega2** Most Recent 2 years, 1 month ago

C obviously

upvoted 1 times

**Marge23** 2 years, 1 month ago

show grants to role; generate results for all granted privileges of the role.

upvoted 1 times

**learner2023** 2 years, 6 months ago

SHOW GRANTS TO ... ROLE role_name Lists all privileges and roles granted to the role.

upvoted 2 times

**KarBiswa** 2 years, 7 months ago

https://docs.snowflake.com/en/sql-reference/sql/show-grants

upvoted 2 times

**Kate0204** 2 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/sql/show-grants.html#show-grants

upvoted 4 times

Question #265 *Topic 1*

Which Query Profile result indicates that a warehouse is sized too small?

A. There are a lot of filter nodes.

B. Bytes are spilling to external storage.

C. The number of processed rows is very high.

D. The number of partitions scanned is the same as partitions total.

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Rob__C** Highly Voted 3 years, 4 months ago

https://community.snowflake.com/s/article/Performance-impact-from-local-and-remote-disk-spilling

upvoted 5 times

**_yyukta** Most Recent 2 years, 2 months ago

B. Bytes are spilling to external storage.

upvoted 1 times

**MultiCloudIronMan** 3 years, 2 months ago

This questions has come before

upvoted 4 times

Question #266 *Topic 1*

What is the default Time Travel retention period?

A. 1 day

B. 7 days

C. 45 days

D. 90 days

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Kate0204** Highly Voted 3 years, 5 months ago

1 is default, but in Enterprise can be extended to max 90 days.

upvoted 8 times

**PrashantGupta1616** Most Recent 1 year, 11 months ago

In Snowflake, the default Time Travel retention period is: A. 1 day By default, Snowflake retains historical data for one day, allowing users to query, clone, and restore data from that period.

upvoted 1 times

**_yyukta** 2 years, 2 months ago

A. 1 day

upvoted 1 times

**Avi0407** 2 years, 6 months ago

Bro why bro why... A - 1 day is the Answer 90 days is maximum... if 90 is defaut company total project budget will go to snowflake maintenance only....

upvoted 1 times

**singhks** 2 years, 8 months ago

1 Day is the default retention time

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**c0d3g** 2 years, 10 months ago

1 day is default.

upvoted 1 times

**learner2023** 3 years ago

No tasks are required to enable Time Travel. It is automatically enabled with the standard, 1-day retention period. However, you may wish to upgrade to Snowflake Enterprise Edition to enable configuring longer data retention periods of up to 90 days for databases, schemas, and tables.

upvoted 2 times

**EmiB** 3 years, 1 month ago

A. 1 day

upvoted 3 times

**miadiesel** 3 years, 4 months ago

default of 1 day https://docs.snowflake.com/en/user-guide/data-time-travel.html#data-retention-period

upvoted 4 times

Question #267 *Topic 1*

Which of the following are best practice recommendations that should be considered when loading data into Snowflake? (Choose two.)

A. Load files that are approximately 25 MB or smaller.

B. Remove all dates and timestamps.

C. Load files that are approximately 100-250 MB (or larger).

D. Avoid using embedded characters such as commas for numeric data types.

E. Remove semi-structured data types.

**Suggested Answer:** *C, D*

*Community vote distribution*

CD (100%)

**_yyukta** Most Recent 2 years, 2 months ago

C and D are correct

upvoted 1 times

**SCube** 2 years, 6 months ago

C & D. Ref : https://docs.snowflake.com/en/user-guide/data-load-considerations-prepare

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 2 times

**KarBiswa** 3 years, 1 month ago

correct answer https://docs.snowflake.com/en/user-guide/data-load-considerations-prepare.html#preparing-delimited-text-files

upvoted 3 times

**fahfouhi94** 3 years, 3 months ago

C & D => correct answer

upvoted 3 times

**sailoo** 3 years, 3 months ago

C,D is correct

upvoted 1 times

**SV1122** 3 years, 4 months ago

Avoid embedded characters, such as commas (e.g. 123,456). https://docs.snowflake.com/en/user-guide/data-load-considerations-prepare.html#preparing-delimited-text-files

upvoted 2 times

**Kate0204** 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/data-load-considerations-prepare.html https://www.snowflake.com/blog/best-practices-for-data-ingestion/

upvoted 4 times

Question #268 *Topic 1*

Which schema has the RESOURCE_MONITORS view?

A. ACCOUNT_USAGE

B. READER_ACCOUNT_USAGE

C. INFORMATION_SCHEMA

D. WAREHOUSE_USAGE_SCHEMA

**Suggested Answer:** *B*

*Community vote distribution*

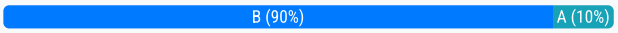


**Kate0204** Highly Voted 3 years, 5 months ago

https://docs.snowflake.com/en/sql-reference/account-usage/resource_monitors.html

upvoted 11 times

**halol** Highly Voted 3 years, 5 months ago

https://docs.snowflake.com/en/sql-reference/account-usage/resource_monitors.html

upvoted 6 times

**ReddySid7** Most Recent 9 months, 1 week ago

ACCOUNT_USAGE: This schema contains the RESOURCE_MONITORS view used to display resource monitors created in your primary account. READER_ACCOUNT_USAGE: This schema also contains a RESOURCE_MONITORS view specifically for tracking monitors in reader accounts managed by the provider. Note for the exam: While the view technically exists in READER_ACCOUNT_USAGE as well, standard exam questions regarding general account resource monitoring always point to ACCOUNT_USAGE as the primary source of truth.

upvoted 1 times

**9b87322** 9 months, 4 weeks ago

Reader_Account_Usage only monitors Reader Account. For your normal account the resource monitor is in ACCOUNT_UAGE. Answer should be A. if it was multiplechoice question both A & B could be been right

upvoted 1 times

**Wzghaer** 10 months, 1 week ago

it can be A ACCOUNT_USAGE, READER_ACCOUNT_USAGE and ORGANIZATION_USAGE the RESOURCE_MONITORS view is available in all of these schemas view

upvoted 1 times

**Altria** 1 year, 1 month ago

-- Access the RESOURCE_MONITORS view SELECT * FROM SNOWFLAKE.READER_ACCOUNT_USAGE.RESOURCE_MONITORS;

upvoted 1 times

**_yyukta** 2 years, 2 months ago

B. READER_ACCOUNT_USAGE

upvoted 1 times

**DavidMiller** 2 years, 4 months ago

B is Correct

upvoted 1 times

**Avi0407** 2 years, 6 months ago

RESOURCE_MONITORS View: This Account Usage view displays the resource monitors that have been created in the reader accounts managed by the account. Note This view is only available in the READER_ACCOUNT_USAGE schema. https://docs.snowflake.com/en/sql-reference/account-usage/resource_monitors

upvoted 3 times

**singhks** 2 years, 8 months ago

B is correct. https://docs.snowflake.com/en/sql-reference/account-usage/resource_monitors#

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**c0d3g** 2 years, 10 months ago

This view is only available in the READER_ACCOUNT_USAGE schema.

upvoted 1 times

**PriyankaTN** 2 years, 10 months ago

answer is B

upvoted 1 times

**abdel88** 2 years, 11 months ago

You can execute this : use database SNOWFLAKE; select * from reader_account_usage.resource_monitors

upvoted 1 times

**isolated_sniper** 3 years ago

RESOURCE_MONITORS View This Account Usage view displays the resource monitors that have been created in the reader accounts managed by the account. Note: This view is only available in the READER_ACCOUNT_USAGE schema. REF: https://docs.snowflake.com/en/sql-reference/account-usage/resource_monitors#

upvoted 2 times

**aroma** 3 years ago

This view is only available in the READER_ACCOUNT_USAGE schema.

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

redaer account usage

upvoted 2 times

Question #269 *Topic 1*

What is the purpose of enabling Federated Authentication on a Snowflake account?

A. Disables the ability to use key pair and basic authentication (e.g., username/password) when connecting

B. Allows dual Multi-Factor Authentication (MFA) when connecting to Snowflake

C. Forces users to connect through a secure network proxy

D. Allows users to connect using secure single sign-on (SSO) through an external identity provider

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

D is correct answer

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. Allows users to connect using secure single sign-on (SSO) through an external identity provider

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**bibisa** 3 years, 6 months ago

In a federated environment, user authentication is separated from user access through the use of one or more external entities that provide independent authentication of user credentials. The authentication is then passed to one or more services, enabling users to access the services through SSO. A federated environment consists of the following components:

upvoted 3 times

**Dan1420** 3 years, 6 months ago

D is correct

upvoted 1 times

**sakis213** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/admin-security-fed-auth.htmlD is correct

upvoted 3 times

**sakis213** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/admin-security-fed-auth.html

upvoted 3 times

**sakis213** 3 years, 10 months ago

D is correct

upvoted 3 times

Question #270 *Topic 1*

Which Snowflake partner category is represented at the top of this diagram (labeled 1)?




A. Business Intelligence

B. Machine Learning and Data Science

C. Security and Governance

D. Data Integration

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Data Integration

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. Data Integration

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**EmiB** 3 years, 7 months ago

D. Data Integration

upvoted 2 times

**SV1122** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/ecosystem-security.html

upvoted 3 times

Question #271 *Topic 1*

Which object types are protected by Fail-safe? (Choose two.)

A. Permanent Tables

B. Temporary Tables

C. External Tables

D. Materialized Views

E. Transient Tables

**Suggested Answer:** *A, D*

*Community vote distribution*

AD (100%)

**SV1122** Highly Voted 3 years, 10 months ago

Snowflake documentation only mentions fail-safe for permanent tables. No mention of materialized views of external tables. Don't know how to pick 2 answers here. https://docs.snowflake.com/en/user-guide/data-failsafe.html

upvoted 6 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Permanent Tables and Materialized views

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A and D are correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**EmiB** 3 years, 7 months ago

Selected Answer: AD

TABLE & MV

upvoted 2 times

**EmiB** 3 years, 7 months ago

TABLE & MVOf course PERMANENT TABLE

upvoted 2 times

**EmiB** 3 years, 7 months ago

Of course PERMANENT TABLE

upvoted 1 times

**fahfouhi94** 3 years, 9 months ago

I think permanent and materlized view

upvoted 2 times

**Naveen90** 3 years, 10 months ago

Answers seem to be correct. Temporary and Transient tables are not protected by fail-safe feature. External table is basically a file, so doesn't make sense for fail-safe.

upvoted 3 times

Question #272 *Topic 1*

Snowflake's approach to the management of system access combines which of the following models? (Choose two.)

A. Security Assertion Markup Language (SAML)

B. Role-Based Access Control (RBAC)

C. Identity Access Management (AM)

D. Create, Read, Update, and Delete (CRUD)

E. Discretionary Access Control (DAC)

F. Mandatory Access Control (MAC)

**Suggested Answer:** *B, E*

Snowflake's approach to access control combines Discretionary Access Control (DAC), where each object has an owner who can grant access to it, and Role-Based Access Control (RBAC), where privileges are granted to roles and roles are granted to users. Reference: https://docs.snowflake.com/en/user-guide/security-access-control-overview

*Community vote distribution*




**Kate0204** Highly Voted 3 years, 11 months ago

DAC + RBAC

upvoted 13 times

**harshagc180** Highly Voted 3 years, 11 months ago

Answer : B & E

upvoted 11 times

**Ashush** Most Recent 1 year, 9 months ago

B & E are correct answers

upvoted 1 times

**VIVEK2025** 1 year, 10 months ago

RBAC+DAC

upvoted 1 times

**Shipra123** 1 year, 11 months ago

B and E are correct answers.

upvoted 1 times

**VLMBL_94** 2 years, 2 months ago

The answer is DAC & RBAC

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B and E are correct

upvoted 1 times

**YoshiSamoZou** 2 years, 4 months ago

DAC + RBAC

upvoted 2 times

**PrashantGupta1616** 2 years, 5 months ago

ANSWER is BE RBAC and DAC

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Role-Based Access Control (RBAC)

upvoted 3 times

**Mike722** 2 years, 8 months ago

B. Role-Based Access Control (RBAC) - Allows other than owners to access/use. E. Discretionary Access Control (DAC) - Allows ownership.

upvoted 2 times

**sandy91** 2 years, 9 months ago

B and E are correct

upvoted 3 times

**azlovemovie** 2 years, 10 months ago

B. Role-Based Access Control (RBAC) E. Discretionary Access Control (DAC)

upvoted 3 times

**Avi0407** 2 years, 12 months ago

Access Control Framework¶ Snowflake's approach to access control combines aspects from both of the following models: Discretionary Access Control (DAC): Each object has an owner, who can in turn grant access to that object. Role-based Access Control (RBAC): Access privileges are assigned to roles, which are in turn assigned to users.

upvoted 2 times

**romega2** 3 years, 1 month ago

E and B

upvoted 2 times

**Amitsnowflake** 3 years, 1 month ago

BE https://docs.snowflake.com/user-guide/security-access-control-overview#label-access-control-overview-roles-system

upvoted 2 times

**Amitsnowflake** 3 years, 1 month ago

BD https://docs.snowflake.com/user-guide/security-access-control-overview#label-access-control-overview-roles-system

upvoted 1 times

Question #273 *Topic 1*

According to Snowflake best practice recommendations, which role should be used to create databases?

A. ACCOUNTADMIN

B. SYSADMIN

C. SECURITYADMIN

D. USERADMIN

**Suggested Answer:** *B*

*Community vote distribution*




**Mallikharjuna452** Most Recent 2 years, 3 months ago

Sysadmin used to create databases and warehouses

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. SYSADMIN

upvoted 1 times

**Afzy** 3 years ago

Chat GPT is saying C While "SECURITYADMIN" is a recommended role for creating databases, it's essential to configure this role's privileges and restrictions appropriately, as well as consider any specific organizational policies and requirements related to database creation and administration.

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**c0d3g** 3 years, 4 months ago

B. SYSADMIN

upvoted 1 times

**OscarWest2023** 3 years, 5 months ago

Yes, sysadmin can create users. Sysadmin is effectively a superuser so it can create multiple objects. However, by least privileges principles, you really shouldn't be using these privileged accounts all the time. There is a reason the useradmin account exists - this is so more users with minimal privs can be created and used. https://docs.snowflake.com/en/user-guide/admin-user-management In the link above, it says "The USERADMIN system role can create users using SQL (CREATE USER). If you prefer to use a custom role for this purpose, grant the CREATE USER privilege on the account to this role.". This suggests to me that snowflake is recommending users be created with useradmin and not sysadmin.The question is: According to Snowflake best practice recommendations, which role should be used to create databases? It is not asking which role should be used to create users?

upvoted 1 times

**OscarWest2023** 3 years, 5 months ago

Selected Answer: D

Yes, sysadmin can create users. Sysadmin is effectively a superuser so it can create multiple objects. However, by least privileges principles, you really shouldn't be using these privileged accounts all the time. There is a reason the useradmin account exists - this is so more users with minimal privs can be created and used.

https://docs.snowflake.com/en/user-guide/admin-user-management

In the link above, it says "The USERADMIN system role can create users using SQL (CREATE USER).

If you prefer to use a custom role for this purpose, grant the CREATE USER privilege on the account to this role.".

This suggests to me that snowflake is recommending users be created with useradmin and not sysadmin.

upvoted 1 times

**c0d3g** 3 years, 4 months ago

The question is: According to Snowflake best practice recommendations, which role should be used to create databases? It is not asking which role should be used to create users?

upvoted 2 times

**SV1122** 3 years, 10 months ago

The system administrator (SYSADMIN) role includes the privileges to create warehouses, databases, and all database objects (schemas, tables, etc.). https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html#:~:text=The%20system%20administrator%20(SYSADMIN)%20role,%2C%20tables%2C%20etc.).

upvoted 4 times

**Kate0204** 3 years, 11 months ago

The system administrator (SYSADMIN) role includes the privileges to create warehouses, databases, and all database objects (schemas, tables, etc.).

upvoted 4 times

Question #274 *Topic 1*

To add or remove search optimization for a table, a user must have which of the following privileges or roles? (Choose two.)

A. The MODIFY privilege on the table

B. The OWNERSHIP privilege on the table

C. A SECURITYADMIN role

D. The ADD SEARCH OPTIMIZATION privilege on the schema that contains the table

E. The SELECT privilege on the table

**Suggested Answer:** *B, D*

*Community vote distribution*

BD (100%)

**Kate0204** Highly Voted 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/search-optimization-service.html#what-access-control-privileges-are-needed-for-the-search-optimization-service

upvoted 7 times

**_yyukta** Most Recent 2 years, 2 months ago

BD are correct

upvoted 1 times

**whiteomax** 2 years, 2 months ago

https://docs.snowflake.com/en/user-guide/search-optimization/enabling

upvoted 1 times

**e_l_i_a_s** 2 years, 3 months ago

correct

upvoted 1 times

**Ices_Eyed** 2 years, 5 months ago

https://docs.snowflake.com/en/user-guide/search-optimization/enabling#required-access-control-privileges

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/search-optimization-service.html#what-access-control-privileges-are-needed-for-the-search-optimization-service

upvoted 2 times

**SV1122** 3 years, 4 months ago

To add, configure, or remove search optimization for a table, you must have the following privileges: You must have OWNERSHIP privilege on the table. You must have ADD SEARCH OPTIMIZATION privilege on the schema that contains the table. https://docs.snowflake.com/en/user-guide/search-optimization-service.html#what-access-control-privileges-are-needed-for-the-search-optimization-service

upvoted 4 times

**nintendogamer64** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/search-optimization-service.html#what-access-control-privileges-are-needed-for-the-search-optimization-service

upvoted 1 times

Question #275 *Topic 1*

While using a COPY command with a Validation_mode parameter, which of the following statements will return an error?

A. Statements that insert a duplicate record during a load

B. Statements that have a specific data type in the source

C. Statements that have duplicate file names

D. Statements that transform data during a load

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**harshagc180** Highly Voted 3 years, 11 months ago

Answer D The VALIDATION_MODE parameter does not support COPY statements that transform data during a load.
https://docs.snowflake.com/en/user-guide/data-load-transform.html#validation-mode-parameter

upvoted 12 times

**necroz** Most Recent 1 year, 11 months ago

VALIDATION_MODE does not support COPY statements that transform data during a load. If the parameter is specified, the COPY statement returns an error.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. Statements that transform data during a load

upvoted 1 times

**Marge23** 3 years, 1 month ago

VALIDATION_MODE Parameter¶ The VALIDATION_MODE parameter does not support COPY statements that transform data during a load.

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**Raju039** 3 years, 7 months ago

Correct D

upvoted 3 times

**AlexbDku** 3 years, 10 months ago

Answer D https://docs.snowflake.com/en/user-guide/data-load-transform.html#validation-mode-parameter

upvoted 3 times

Question #276 *Topic 1*

When is the result set cache no longer available? (Choose two.)

A. When another warehouse is used to execute the query

B. When another user executes the query

C. When the underlying data has changed

D. When the warehouse used to execute the query is suspended

E. When it has been 24 hours since the last query

**Suggested Answer:** *C*

*Community vote distribution*




**Kate0204** Highly Voted 3 years, 11 months ago

C and E

upvoted 24 times

**Ashush** Most Recent 1 year, 9 months ago

C and E are correct answers

upvoted 2 times

**JaQuanda** 1 year, 11 months ago

C + E are both correct

upvoted 1 times

**Shipra123** 1 year, 11 months ago

C and E are correct answers

upvoted 2 times

**Mallikharjuna452** 2 years, 3 months ago

C AND E are correct

upvoted 2 times

**PrashantGupta1616** 2 years, 5 months ago

C and E

upvoted 1 times

**Mitchdu** 2 years, 7 months ago

C and E

upvoted 2 times

**_yyukta** 2 years, 8 months ago

C. When the underlying data has changed

upvoted 1 times

**sandy91** 2 years, 9 months ago

C and E are correct

upvoted 3 times

**JM192** 2 years, 10 months ago

C and E

upvoted 2 times

**azlovemovie** 2 years, 10 months ago

C and E. I'm trying to make both voting 50% each

upvoted 1 times

**GAUURAAV** 2 years, 11 months ago

C and E

upvoted 1 times

**romega2** 3 years, 1 month ago

C and E

upvoted 1 times

**Amitsnowflake** 3 years, 1 month ago

C and E Query result persist only for 24 hours.

upvoted 1 times

**Marge23** 3 years, 1 month ago

C & E correct, if the use_cached_result is set to extended period.

upvoted 1 times

**Taraka9** 3 years, 2 months ago

CE is correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

CE Correct

upvoted 1 times

Question #277 *Topic 1*

What is the recommended file sizing for data loading using Snowpipe?

A. A compressed file size greater than 100 MB, and up to 250 MB

B. A compressed file size greater than 100 GB, and up to 250 GB

C. A compressed file size greater than 10 MB, and up to 100 MB

D. A compressed file size greater than 1 GB, and up to 2 GB

**Suggested Answer:** *A*

*Community vote distribution*

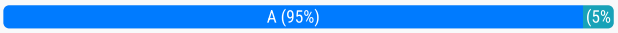


**Kate0204** Highly Voted 3 years, 11 months ago

Loading data files roughly 100-250 MB in size or larger reduces the overhead charge relative to the amount of total data loaded to the point where the overhead cost is immaterial.

upvoted 8 times

**Kate0204** Highly Voted 4 months ago

Selected Answer: D, E, F

Snowflake is provided as Software-as-a-Service (SaaS) that runs completely on cloud infrastructure. This means that all three layers of Snowflake's architecture (storage, compute, and cloud services) are deployed and managed entirely on a selected cloud platform. In addition, Snowflake's virtual warehouses and cloud services layers are similarly deployed across three availability zones in a region.

upvoted 8 times

**Mallikharjuna452** 2 years, 3 months ago

A is correct Answer

upvoted 1 times

**tpiscaglia** 4 months ago

Selected Answer: D, E, F

D,E,F https://docs.snowflake.com/en/user-guide/intro-cloud-platforms https://developers.snowflake.com/wp-content/uploads/2021/06/Snowflake-High-Availability-for-Data-Apps-Whitepaper.pdf

upvoted 4 times

**_yyukta** 2 years, 8 months ago

C. When the underlying data has changed

upvoted 1 times

**EmiB** 4 months ago

Selected Answer: D, E, F

D,E & F - correct

upvoted 3 times

**Afzy** 3 years ago

i think A is the answer

upvoted 1 times

**fahfouhi94** 4 months ago

Selected Answer: D, E, F

D,E & F are true

upvoted 2 times

**Marge23** 3 years, 1 month ago

correct

upvoted 1 times

**KarBiswa** Most Recent 4 months ago

any supportive document?

upvoted 1 times

**singhks** 3 years, 2 months ago

A is correct

upvoted 1 times

**c0d3g** 3 years, 4 months ago

There is a fixed, per-file overhead charge for Snowpipe in addition to the compute processing costs. We recommend files at least above 10 MB on average, with files in the 100 to 250 MB range offering the best cost-to-performance ratio. https://www.snowflake.com/blog/best-practices-for-data-ingestion/#:~:text=Recommended%20file%20size%20for%20Snowpipe%20and%20cost%20considerations&text=We%20recommend%20files%20at%20least,cost%2Dto%2Dperformance%20ratio.

upvoted 2 times

**tstander** 3 years, 4 months ago

I suspect these "best practices" where updated with time. I have documentation dated June 2020 from attending a Snowflake Quickstart Training (by Snowflake Proffessional Services) that aligns with the moderated answer.

upvoted 1 times

**tpiscaglia** 3 years, 7 months ago

A https://www.snowflake.com/blog/best-practices-for-data-ingestion/#:~:text=Recommended%20file%20size%20for%20Snowpipe%20and%20cost%20considerations&text=We%20recommend%20files%20at%20least,cost%2Dto%2Dperformance%20ratio.

upvoted 2 times

**EmiB** 3 years, 7 months ago

Loading data files roughly 100-250 MB in size or larger reduces the overhead charge relative to the amount of total data loaded to the point where the overhead cost is immaterial.

upvoted 1 times

**ojoa** 3 years, 11 months ago

Loading data files roughly 100-250 MB in size or larger reduces the overhead charge relative to the amount of total data loaded to the point where the overhead cost is immaterial.

upvoted 3 times

**ojoa** 3 years, 11 months ago

Loading data files roughly 100-250 MB in size or larger reduces the overhead charge relative to the amount of total data loaded to the point where the overhead cost is immaterial.MB not GB

upvoted 1 times

**EmiB** 3 years, 7 months ago

MB not GB

upvoted 1 times

Question #278 *Topic 1*

Which statements are true concerning Snowflake's underlying cloud infrastructure? (Choose three.)

A. Snowflake data and services are deployed in a single availability zone within a cloud provider's region.

B. Snowflake data and services are available in a single cloud provider and a single region; the use of multiple cloud providers is not supported.

C. Snowflake can be deployed in a customer's private cloud using the customer's own compute and storage resources for Snowflake compute and storage.

D. Snowflake uses the core compute and storage services of each cloud provider for its own compute and storage.

E. All three layers of Snowflake's architecture (storage, compute, and cloud services) are deployed and managed entirely on a selected cloud platform.

F. Snowflake data and services are deployed in at least three availability zones within a cloud provider's region.

**Suggested Answer:** *D, E, F*

*Community vote distribution*

DEF (100%)

**Kate0204** Highly Voted 3 years, 11 months ago

Selected Answer: DEF

Snowflake is provided as Software-as-a-Service (SaaS) that runs completely on cloud infrastructure. This means that all three layers of Snowflake's architecture (storage, compute, and cloud services) are deployed and managed entirely on a selected cloud platform.

In addition, Snowflake's virtual warehouses and cloud services layers are similarly deployed across three availability zones in a region.

upvoted 8 times

**Kate0204** Highly Voted 3 years, 11 months ago

Snowflake is provided as Software-as-a-Service (SaaS) that runs completely on cloud infrastructure. This means that all three layers of Snowflake's architecture (storage, compute, and cloud services) are deployed and managed entirely on a selected cloud platform. In addition, Snowflake's virtual warehouses and cloud services layers are similarly deployed across three availability zones in a region.any supportive document?

upvoted 8 times

**KarBiswa** 3 years, 8 months ago

any supportive document?

upvoted 1 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

d,e,f are correct options

upvoted 1 times

**_yyukta** 2 years, 8 months ago

DEF are correct

upvoted 1 times

**0e504b5** 2 years, 9 months ago

Question: why is B not correct?I would assume it's because each account in an organization can be deployed in a separate cloud service provider https://docs.snowflake.com/en/user-guide/intro-cloud-platforms On each platform, Snowflake provides one or more regions where the account is provisioned. If your organization's other cloud services are already hosted on one of these platforms, you can choose to host all your Snowflake accounts on the same platform. However, you can also choose to host your accounts on a different platform.

upvoted 1 times

**0e504b5** 2 years, 9 months ago

Question: why is B not correct?

upvoted 1 times

**ahadh7621** 2 years, 5 months ago

I would assume it's because each account in an organization can be deployed in a separate cloud service provider https://docs.snowflake.com/en/user-guide/intro-cloud-platforms On each platform, Snowflake provides one or more regions where the account is provisioned. If your organization's other cloud services are already hosted on one of these platforms, you can choose to host all your Snowflake accounts on the same platform. However, you can also choose to host your accounts on a different platform.

upvoted 1 times

**Needium** 3 years, 6 months ago

D,E & F are true

upvoted 1 times

**tpiscaglia** 3 years, 7 months ago

D,E,F https://docs.snowflake.com/en/user-guide/intro-cloud-platforms https://developers.snowflake.com/wp-content/uploads/2021/06/Snowflake-High-Availability-for-Data-Apps-Whitepaper.pdf

upvoted 4 times

**EmiB** 3 years, 7 months ago

D,E & F - correct

upvoted 3 times

**KarBiswa** 3 years, 8 months ago

DE is seems correct but on F is there any supported document ?

upvoted 1 times

**KarBiswa** 3 years, 8 months ago

DE is seems correct but on F is there any supported document ?This whitepaper has the following: The Snowflake deployment extends across three availability zones. https://developers.snowflake.com/wp-content/uploads/2021/06/Snowflake-High-Availability-for-Data-Apps-Whitepaper.pdf

upvoted 1 times

**tpiscaglia** 3 years, 7 months ago

This whitepaper has the following: The Snowflake deployment extends across three availability zones. https://developers.snowflake.com/wp-content/uploads/2021/06/Snowflake-High-Availability-for-Data-Apps-Whitepaper.pdf

upvoted 1 times

**fahfouhi94** 3 years, 9 months ago

D,E & F are true

upvoted 2 times

Question #279 *Topic 1*

A user unloaded a Snowflake table called mytable to an internal stage called mystage.

Which command can be used to view the list of files that has been uploaded to the stage?

A. list @mytable;

B. list @%mytable;

C. list @%mystage;

D. list @mystage;

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**halol** Highly Voted 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/sql/list.html

upvoted 5 times

**Rajivnb** Highly Voted 3 years ago

User Stage should start with @~ Table Stage should start with @% Named Stage or Internal Stage of Snowflake is just @.

upvoted 5 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

list @mystage

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. list @mystage;

upvoted 1 times

**EmiB** 3 years, 7 months ago

B. list @%mytable; -- copy into @%mytable FROM mytable; D. list @mystage; -- copy into @mystage FROM mytable;

upvoted 2 times

**fahfouhi94** 3 years, 9 months ago

correct answer

upvoted 2 times

**harshagc180** 3 years, 11 months ago

Correct Answer D

upvoted 2 times

Question #280 *Topic 1*

What is a best practice after creating a custom role?

A. Create the custom role using the SYSADMIN role.

B. Assign the custom role to the SYSADMIN role.

C. Assign the custom role to the PUBLIC role.

D. Add _CUSTOM to all custom role names.

**Suggested Answer:** *B*

*Community vote distribution*

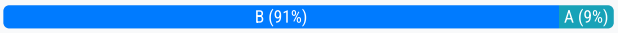


**SnowProCertDec22** Highly Voted 3 years, 10 months ago

answer is B Custom roles (i.e. any roles other than the system-defined roles) can be created by the USERADMIN role (or a higher role) as well as by any role to which the CREATE ROLE privilege has been granted. By default, a newly-created role is not assigned to any user, nor granted to any other role. When creating roles that will serve as the owners of securable objects in the system, Snowflake recommends creating a hierarchy of custom roles, with the top-most custom role assigned to the system role SYSADMIN.

upvoted 9 times

**visionary4ever** Most Recent 1 year, 2 months ago

SYSADMIN creates custom roles Granting custom roles back to SYSADMIN → Not possible, that's backwards.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B Represents correct answer

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Assign the custom role to the SYSADMIN role.

upvoted 1 times

**Marge23** 3 years, 1 month ago

SYSADMIN (aka System Administrator) Role that has privileges to create warehouses and databases (and other objects) in an account. If, as recommended, you create a role hierarchy that ultimately assigns all custom roles to the SYSADMIN role, this role also has the ability to grant privileges on warehouses, databases, and other objects to other roles.

upvoted 1 times

**singhks** 3 years, 2 months ago

B is the right answer

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**EmiB** 3 years, 7 months ago

B. Assign the custom role to the SYSADMIN role.

upvoted 2 times

**miadiesel** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#role-hierarchy-and-privilege-inheritance
https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#role-hierarchy-and-privilege-inheritance

upvoted 4 times

**harshagc180** 3 years, 11 months ago

Answer A https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#Sorry, correct answer is B

upvoted 1 times

**harshagc180** 3 years, 11 months ago

Answer A https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#

upvoted 1 times

**harshagc180** 3 years, 10 months ago

Sorry, correct answer is B

upvoted 1 times

Question #281 *Topic 1*

Which is the MINIMUM required Snowflake edition that a user must have if they want to use AWS/Azure Privatelink or Google Cloud Private Service Connect?

A. Standard

B. Premium

C. Enterprise

D. Business Critical

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**arpit_dataguy** Highly Voted 3 years, 10 months ago

d https://docs.snowflake.com/en/user-guide/admin-security-privatelink.html

upvoted 12 times

**miadiesel** Highly Voted 3 years, 10 months ago

This feature requires Business Critical (or higher). https://docs.snowflake.com/en/user-guide/admin-security-privatelink.html

upvoted 8 times

**PrasantSadangi** Most Recent 1 year, 7 months ago

requires Business Critical (or higher).

upvoted 1 times

**KIYREY** 1 year, 10 months ago

BUSINESS CRITICAL

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

Business critical

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. Business Critical

upvoted 1 times

**VASI250** 2 years, 10 months ago

This feature requires Business Critical (or higher) as mentioned in documentation

upvoted 1 times

**Marge23** 3 years, 1 month ago

correct

upvoted 1 times

**singhks** 3 years, 2 months ago

Business Critical or higher supports Private Link

upvoted 1 times

**OscarWest2023** 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/intro-editions

upvoted 1 times

**aroma** 3 years, 6 months ago

AWS PrivateLink & Snowflake BUSINESS CRITICAL FEATURE This feature requires Business Critical (or higher).

upvoted 1 times

**hrushi98** 3 years, 7 months ago

Enterprise and use PrivateLink but with an additional fee

upvoted 2 times

**EmiB** 3 years, 7 months ago

D. Business Critical

upvoted 1 times

**miadiesel** 3 years, 9 months ago

Private connectivity to the Snowflake service requires Business Critical (or higher). https://docs.snowflake.com/en/user-guide/private-snowflake-service.html#private-connectivity-to-the-snowflake-service

upvoted 2 times

**SnowProCertDec22** 3 years, 10 months ago

Snowflake logo in black (no text) Business Critical Feature This feature requires Business Critical (or higher).

upvoted 2 times

Question #282 *Topic 1*

Which of the following query profiler variables will indicate that a virtual warehouse is not sized correctly for the query being executed?

A. Bytes sent over the network

B. Synchronization

C. Initialization

D. Remote spillage

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**n21007** Highly Voted 3 years, 8 months ago

D is correct

upvoted 5 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Remote spilling

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D. Remote spillage

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**girgir** 3 years, 8 months ago

For some operations (e.g. duplicate elimination for a huge data set), the amount of memory available for the compute resources used to execute the operation might not be sufficient to hold intermediate results. As a result, the query processing engine will start spilling the data to local disk. If the local disk space is not sufficient, the spilled data is then saved to remote disks. https://docs.snowflake.com/en/user-guide/ui-query-profile

upvoted 2 times

Question #283 *Topic 1*

Which of the following Snowflake capabilities are available in all Snowflake editions? (Choose two.)

A. Customer-managed encryption keys through Tri-Secret Secure

B. Automatic encryption of all data

C. Up to 90 days of data recovery through Time Travel

D. Object-level access control

E. Column-level security to apply data masking policies to tables and views

**Suggested Answer:** *B*

*Community vote distribution*




**Kate0204** Highly Voted 3 years, 11 months ago

Busines Critical and VPS A. Customer-managed encryption keys through Tri-Secret Secure Enterprise: C. Up to 90 days of data recovery through Time Travel E. Column-level security to apply data masking policies to tables and views All: B. Automatic encryption of all data D. Object-level access control

upvoted 45 times

**eldorado92** Highly Voted 3 years, 10 months ago

B&D is the answer

upvoted 13 times

**Ashush** Most Recent 1 year, 9 months ago

B & D are correct answers

upvoted 1 times

**idris_qim** 1 year, 9 months ago

B&D is the answer

upvoted 1 times

**d22770a** 2 years ago

Ans is B and D

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

B and D is the answer

upvoted 1 times

**PrashantGupta1616** 2 years, 5 months ago

B&D is the answer

upvoted 1 times

**jerintom** 2 years, 5 months ago

Option B and D

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Automatic encryption of all data

upvoted 2 times

**sandy91** 2 years, 9 months ago

B and D are correct

upvoted 1 times

**GAUURAAV** 2 years, 11 months ago

B and D

upvoted 1 times

**SCube** 3 years ago

B & D is the answer . Ref : https://docs.snowflake.com/en/user-guide/intro-editions

upvoted 1 times

**Taraka9** 3 years, 2 months ago

B and D are correct

upvoted 1 times

**conrad2023** 3 years, 2 months ago

B & D are correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

BD Correct

upvoted 1 times

**c0d3g** 3 years, 4 months ago

B AND D

upvoted 1 times

**RajuNaik** 3 years, 6 months ago

BD are correct

upvoted 1 times

Question #284 *Topic 1*

A PUT command can be used to stage local files from which Snowflake interface?

A. SnowSQL

B. Snowflake classic web interface (UI)

C. Snowsight

D. .NET driver

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Snowsql

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A. SnowSQL

upvoted 1 times

**0e504b5** 2 years, 9 months ago

https://docs.snowflake.com/en/sql-reference/sql/put The command cannot be executed from the Worksheets Worksheet tab page in either Snowflake web interface; instead, use the SnowSQL client or Drivers to upload data files, or check the documentation for a specific Snowflake client to verify support for this command.

upvoted 1 times

**Marge23** 3 years, 1 month ago

snowsql

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**fahfouhi94** 3 years, 9 months ago

correct answer

upvoted 3 times

**leozhang** 3 years, 9 months ago

PUT command Usage The command cannot be executed from the Worksheets page in the Snowflake web interface; instead, use the SnowSQL client to upload data files, or check the documentation for the specific Snowflake client to verify support for this command.
https://roboquery.com/app/syntax-put-command-snowflake

upvoted 4 times

Question #285 *Topic 1*

Which of the following indicates that it may be appropriate to use a clustering key for a table? (Choose two.)

A. The table contains a column that has very low cardinality.

B. DML statements that are being issued against the table are blocked.

C. The table has a small number of micro-partitions.

D. Queries on the table are running slower than expected.

E. The clustering depth for the table is large.

**Suggested Answer:** *D, E*

*Community vote distribution*

DE (100%)

**halol** Highly Voted 3 years, 10 months ago

D&E https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html#:~:text=for%20a%20Table-,What,-is%20a%20Clustering

upvoted 9 times

**Kate0204** Highly Voted 3 years, 11 months ago

DE I think

upvoted 6 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

D and E are correc answers

upvoted 1 times

**_yyukta** 2 years, 8 months ago

DE are correct

upvoted 1 times

**e_l_i_a_s** 2 years, 9 months ago

Some general indicators that can help determine whether to define a clustering key for a table include: Queries on the table are running slower than expected or have noticeably degraded over time. The clustering depth for the table is large.

upvoted 1 times

**Marge23** 3 years, 1 month ago

correct ans

upvoted 1 times

**v1tjkee** 3 years, 1 month ago

Why not A?

upvoted 1 times

**v1tjkee** 3 years, 1 month ago

Why not A?https://docs.snowflake.com/en/user-guide/tables-clustering-keys A column with very low cardinality might yield only minimal pruning, such as a column named IS_NEW_CUSTOMER that contains only Boolean values. In other words, very low cardinality columns may not be worth clustering (the way I read it).

upvoted 1 times

**0e504b5** 2 years, 9 months ago

https://docs.snowflake.com/en/user-guide/tables-clustering-keys A column with very low cardinality might yield only minimal pruning, such as a column named IS_NEW_CUSTOMER that contains only Boolean values. In other words, very low cardinality columns may not be worth clustering (the way I read it).

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**girgir** 3 years, 8 months ago

DE Queries on the table are running slower than expected or have noticeably degraded over time. The clustering depth for the table is large.
https://docs.snowflake.com/en/user-guide/tables-clustering-keys#label-considerations-for-choosing-clustering

upvoted 2 times

**fahfouhi94** 3 years, 9 months ago

I think D E

upvoted 2 times

**AlexbDku** 3 years, 10 months ago

Answer is D and E https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html#what-is-a-clustering-key

upvoted 5 times

Question #286 *Topic 1*

Which cache type is used to cache data output from SQL queries?

A. Metadata cache

B. Result cache

C. Remote cache

D. Local file cache

**Suggested Answer:** *B*

*Community vote distribution*

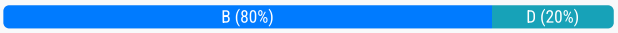


**Kate0204** Highly Voted 3 years, 11 months ago

Result Cache: Which holds the results of every query executed in the past 24 hours. These are available across virtual warehouses, so query results returned to one user is available to any other user on the system who executes the same query, provided the underlying data has not changed. Local Disk Cache: Which is used to cache data used by SQL queries. Whenever data is needed for a given query it's retrieved from the Remote Disk storage, and cached in SSD and memory.

upvoted 11 times

**Kate0204** Highly Voted 3 years, 11 months ago

Result Cache: Which holds the results of every query executed in the past 24 hours. These are available across virtual warehouses, so query results returned to one user is available to any other user on the system who executes the same query, provided the underlying data has not changed. Local Disk Cache: Which is used to cache data used by SQL queries. Whenever data is needed for a given query it's retrieved from the Remote Disk storage, and cached in SSD and memory.Result Query is stored at Cloud Service layer

upvoted 11 times

**shyemko** 3 years, 7 months ago

Result Query is stored at Cloud Service layer

upvoted 2 times

**37ceea2** Most Recent 1 year, 9 months ago

I am looking what Kate 02024 wrote. I am seeing that the question is about SQL queries, in this case, answer is D.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

Results cache

upvoted 2 times

**_yyukta** 2 years, 8 months ago

B. Result cache

upvoted 1 times

**Marge23** 3 years, 1 month ago

result cache

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**EmiB** 3 years, 8 months ago

B is correct

upvoted 1 times

Question #287 *Topic 1*

Which of the following describes how clustering keys work in Snowflake?

A. Clustering keys update the micro-partitions in place with a full sort, and impact the DML operations.

B. Clustering keys sort the designated columns over time, without blocking DML operations.

C. Clustering keys create a distributed, parallel data structure of pointers to a table's rows and columns.

D. Clustering keys establish a hashed key on each node of a virtual warehouse to optimize joins at run-time.

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

B. Clustering keys sort the designated columns over time, without blocking DML operations.

upvoted 1 times

**ahadh7621** 2 years, 5 months ago

B. https://docs.snowflake.com/en/user-guide/tables-auto-reclustering Non-blocking DML Automatic Clustering is transparent and does not block DML statements issued against tables while they are being reclustered.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Clustering keys sort the designated columns over time, without blocking DML operations.

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

correct

upvoted 1 times

**tpiscaglia** 3 years, 7 months ago

B https://docs.snowflake.com/en/user-guide/tables-clustering-keys#what-is-a-clustering-key

upvoted 1 times

**knagnayak** 3 years, 8 months ago

correct answer is B

upvoted 2 times

**fahfouhi94** 3 years, 9 months ago

correct answer is B

upvoted 1 times

**harshagc180** 3 years, 11 months ago

Correct Answer B

upvoted 4 times

Question #288 *Topic 1*

Which of the following operations require the use of a running virtual warehouse? (Choose two.)

A. Downloading data from an internal stage

B. Listing files in a stage

C. Executing a stored procedure

D. Altering a table

E. Querying data from a materialized view

**Suggested Answer:** *C, E*

*Community vote distribution*




**SnowProCertDec22** Highly Voted 3 years, 10 months ago

C & E https://docs.snowflake.com/en/user-guide/warehouses.html

upvoted 14 times

**MultiCloudIronMan** 1 year, 8 months ago

Selected Answer: CD

Material view does not require a warehouse

upvoted 1 times

**MultiCloudIronMan** 1 year, 8 months ago

Material view does not require a warehouseSorry I think virtual warehouse is needed for material view but it fetches the data quicker than running a full table scan.

upvoted 1 times

**MultiCloudIronMan** Most Recent 1 year, 8 months ago

Sorry I think virtual warehouse is needed for material view but it fetches the data quicker than running a full table scan.

upvoted 1 times

**Mallikharjuna452** 2 years, 3 months ago

C and E are correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

CE are correct

upvoted 1 times

**Afzy** 3 years ago

C and E is the answer

upvoted 1 times

**Marge23** 3 years, 1 month ago

Performing DML operations, such as: Updating rows in tables (DELETE , INSERT , UPDATE). Loading data into tables (COPY INTO ). Unloading data from tables (COPY INTO ) https://docs.snowflake.com/en/user-guide/warehouses

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**c0d3g** 3 years, 4 months ago

These 3 do not require active Warehouse: Downloading data from an internal stage Listing files in a stage Creating and Altering table.

upvoted 3 times

**tarekff** 3 years, 5 months ago

Performing DML operations, such as: Updating rows in tables (DELETE , INSERT , UPDATE). (Answer D) Loading data into tables (COPY INTO ). Unloading data from tables (COPY INTO ). (Answer A) a warehouse must be running and in use for the session.Cannot be D. Altering table does not require a compute. Alter table is not altering data. its altering the definition. like changing Column type, removing a column or adding a new column. basically CREATE, ALTER. not the DML. A also cannot be correct. They said downloading data. it did not say clearly that Loading/Unloading data. downloading could be a downloading a file from internal stage to local machine.

upvoted 2 times

**tarekff** 3 years, 5 months ago

**Selected Answer: AD**

Performing DML operations, such as: Updating rows in tables (DELETE , INSERT , UPDATE). (Answer D) Loading data into tables (COPY INTO ). Unloading data from tables (COPY INTO ). (Answer A) a warehouse must be running and in use for the session.

upvoted 2 times

**Rajivnb** 3 years ago

Cannot be D. Altering table does not require a compute. Alter table is not altering data. its altering the definition. like changing Column type, removing a column or adding a new column. basically CREATE, ALTER. not the DML. A also cannot be correct. They said downloading data. it did not say clearly that Loading/Unloading data. downloading could be a downloading a file from internal stage to local machine.

upvoted 1 times

**vini15** 3 years, 6 months ago

Answer is correct C&D Materialised views are pre computed it does require warehouse.

upvoted 3 times

**Omkarshingare** 3 years, 7 months ago

C & D IS CORRECT

upvoted 1 times

**N89** 3 years, 8 months ago

Not sure if materialized view will require a running cluster as data is already computed and may reside in cache. But alter table will definitely require a running cluster.

upvoted 1 times

**Kate0204** 3 years, 11 months ago

Hmmm not sure about this, maybe E

upvoted 2 times

Question #289 *Topic 1*

What is used to limit the credit usage of a virtual warehouse within a Snowflake account?

A. Load monitor

B. Resource monitor

C. Query Profile

D. Stream

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Mallikharjuna452** Most Recent 2 years, 3 months ago

Resource monitor

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B. Resource monitor

upvoted 1 times

**Marge23** 3 years, 1 month ago

RM correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**fahfouhi94** 3 years, 9 months ago

ressource monitor of course

upvoted 2 times

**SnowProCertDec22** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/resource-monitors.html#:~:text=A%20resource%20monitor%20can%20be,and%20how%20long%20it%20runs.

upvoted 3 times

Question #290 *Topic 1*

What are the benefits of the replication feature in Snowflake? (Choose two.)

A. Disaster recovery

B. Time Travel

C. Fail-safe

D. Database failover and failback

E. Data security

**Suggested Answer:** *A, D*

*Community vote distribution*

AD (100%)

**MultiCloudIronMan** Highly Voted 3 years, 8 months ago

Replication is not a security model infact it increases risk exposure, its for business continuity.E is also correct. https://www.snowflake.com/trending/what-is-data-replication#:~:text=DATABASE%20REPLICATION%20SOLUTIONS&text=Customers%20using%20Snowflake%20Standard%20edition,portability%20to%20facilitate%20account%20migrations

upvoted 8 times

**MultiCloudIronMan** Highly Voted 3 years, 8 months ago

Selected Answer: AD

Replication is not a security model infact it increases risk exposure, its for business continuity.

upvoted 8 times

**MultiCloudIronMan** 3 years, 3 months ago

E is also correct. https://www.snowflake.com/trending/what-is-data-replication#:~:text=DATABASE%20REPLICATION%20SOLUTIONS&text=Customers%20using%20Snowflake%20Standard%20edition,portability%20to%20facilitate%20account%20migrations

upvoted 1 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A and D are correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

AD are correct

upvoted 1 times

**Marge23** 3 years, 1 month ago

its for business continuity.

upvoted 1 times

**AS314** 3 years, 9 months ago

https://www.snowflake.com/trending/what-is-data-replication#:~:text=DATABASE%20REPLICATION%20SOLUTIONS&text=Customers%20using%20Snowflake%20Standard%20edition,portability%20to%20facilitate%20account%20migrations so I think D and E are correct

upvoted 3 times

Question #291 *Topic 1*

Which of the following roles are recommended to create and manage users and roles? (Choose two.)

A. SYSADMIN

B. SECURITYADMIN

C. PUBLIC

D. ACCOUNTADMIN

E. USERADMIN

**Suggested Answer:** *B, E*

*Community vote distribution*




**harshagc180** Highly Voted 3 years, 5 months ago

Correct Answer : SECURITYADMIN and USERADMIN https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#roles

upvoted 17 times

**yaho5** Most Recent 2 years, 1 month ago

First - USERADMIN, then SECURITYADMIN

upvoted 1 times

**_yyukta** 2 years, 2 months ago

BE are correct

upvoted 1 times

**azlovemovie** 2 years, 4 months ago

B and E are correct.

upvoted 1 times

**Pradeep_7869** 2 years, 6 months ago

B and E

upvoted 1 times

**Marge23** 2 years, 7 months ago

SECURITYADMIN (aka Security Administrator) Role that can manage any object grant globally, as well as create, monitor, and manage users and roles. More specifically, this role: Is granted the MANAGE GRANTS security privilege to be able to modify any grant, including revoking it. Inherits the privileges of the USERADMIN role via the system role hierarchy (i.e. USERADMIN role is granted to SECURITYADMIN). USERADMIN (aka User and Role Administrator) Role that is dedicated to user and role management only. More specifically, this role: Is granted the CREATE USER and CREATE ROLE security privileges. Can create users and roles in the account. This role can also manage users and roles that it owns. Only the role with the OWNERSHIP privilege on an object (i.e. user or role), or a higher role, can modify the object properties.

upvoted 3 times

**timcooks** 2 years, 7 months ago

SYSTEMADMIN AND SECURITYADMIN

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**c0d3g** 2 years, 10 months ago

SECURITYADMIN inherits from USERADMIN

upvoted 1 times

**ShagunMittal** 3 years, 1 month ago

SECURITYADMIN and USERADMIN

upvoted 2 times

**ntechkuma** 3 years, 1 month ago

sysadmin and securityadmin

upvoted 1 times

**EmiB** 3 years, 2 months ago

E - USERADMIN (as SECURITYADMIN inherits from USERADMIN) or EB - USERADMIN & SECURITYADMIN

upvoted 1 times

**muks86** 3 years, 2 months ago

USER ADmin and security admin

upvoted 1 times

**subbie** 3 years, 4 months ago

B&E are correct answer.

upvoted 4 times

**Kate0204** 3 years, 5 months ago

What about USERADMIN?

upvoted 3 times

Question #292 *Topic 1*

When can a newly configured virtual warehouse start running SQL queries?

A. After 50% of the warehouse provisioning has completed

B. During the time slots defined by the ACCOUNTADMIN

C. When the warehouse provisioning is completed

D. After the warehouse replication is completed

**Suggested Answer:** *C*

*Community vote distribution*

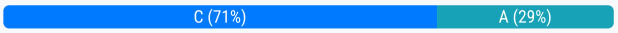


**_yyukta** 2 years, 2 months ago

C is correct

upvoted 1 times

**0e504b5** 2 years, 3 months ago

C is correct

upvoted 1 times

**romega2** 2 years, 7 months ago

C is my answer

upvoted 2 times

**Marge23** 2 years, 7 months ago

correct

upvoted 1 times

**chaitanya_132** 2 years, 9 months ago

BOTH R CORRECT first we preferred to go option d like complete provisioned

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**Dalo** 2 years, 11 months ago

the point is "newly" configured virtual warehouse. so answer is C

upvoted 1 times

**isolated_sniper** 3 years ago

Answer - C. Snowflake does not begin executing SQL statements submitted to a warehouse until all of the compute resources for the warehouse are successfully provisioned, unless any of the resources fail to provision. Ref: https://docs.snowflake.com/en/user-guide/warehouses-tasks# A would be correct for a warehouse that is repaired.

upvoted 3 times

**learner2023** 3 years ago

During the repair process, the warehouse starts processing SQL statements once 50% or more of the requested compute resources are successfully provisioned.

upvoted 1 times

**EmiB** 3 years, 2 months ago

**Selected Answer: A**

It's a tricky question - both A&C could be correct. Snowflake does not begin executing SQL statements submitted to a warehouse until all of the compute resources for the warehouse are successfully provisioned, unless any of the resources fail to provision: - If any of the compute resources for the warehouse fail to provision during start-up, Snowflake attempts to repair the failed resources. - During the repair process, the warehouse starts processing SQL statements once 50% or more of the requested compute resources are successfully provisioned.

upvoted 3 times

**EmiB** 3 years, 2 months ago

It's a tricky question - both A&C could be correct. Snowflake does not begin executing SQL statements submitted to a warehouse until all of the compute resources for the warehouse are successfully provisioned, unless any of the resources fail to provision: - If any of the compute resources for the warehouse fail to provision during start-up, Snowflake attempts to repair the failed resources. - During the repair process, the warehouse starts processing SQL statements once 50% or more of the requested compute resources are successfully provisioned.In this question, it states a NEWLY configured Virtual warehouse. It is C only

upvoted 3 times

**ShawnZh** Most Recent 1 year, 1 month ago

In this question, it states a NEWLY configured Virtual warehouse. It is C only

upvoted 1 times

**BobCui** 3 years, 3 months ago

C is correct

upvoted 2 times

Question #293 *Topic 1*

What actions will prevent leveraging of the ResultSet cache?

A. Removing a column from the query SELECT list

B. Stopping the virtual warehouse that the query is running against

C. If the result has not been reused within the last 12 hours

D. Executing the RESULTS_SCAN() table function

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**_yyukta** Most Recent 2 years, 2 months ago

A is correct

upvoted 1 times

**Marge23** 2 years, 7 months ago

exactly maches previous query.

upvoted 2 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**shyemko** 3 years, 1 month ago

A is correct

upvoted 1 times

**EmiB** 3 years, 2 months ago

A & D both are correct

upvoted 1 times

**EmiB** 3 years, 2 months ago

A & D both are correctOnly A, sorry. D will force the use of Result Cache.

upvoted 1 times

**EmiB** 3 years, 2 months ago

Only A, sorry. D will force the use of Result Cache.

upvoted 2 times

**Yogesh1984** 3 years, 2 months ago

A & D both are correct

upvoted 1 times

Question #294 *Topic 1*

Which of the following are benefits of micro-partitioning? (Choose two.)

A. Micro-partitions cannot overlap in their range of values.

B. Micro-partitions are immutable objects that support the use of Time Travel.

C. Micro-partitions can reduce the amount of I/O from object storage to virtual warehouses.

D. Rows are automatically stored in sorted order within micro-partitions.

E. Micro-partitions can be defined on a schema-by-schema basis.

**Suggested Answer:** *B, C*

*Community vote distribution*

BC (100%)

**arpit_dataguy** Highly Voted 3 years, 4 months ago

Answer -- BC

upvoted 11 times

**BigDataBB** Highly Voted 3 years, 3 months ago

As wrote here: https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions.html#benefits-of-micro-partitioning - In contrast to traditional static partitioning, Snowflake micro-partitions are derived automatically; they don't need to be explicitly defined up-front or maintained by users. - As the name suggests, micro-partitions are small in size (50 to 500 MB, before compression), which enables extremely efficient DML and fine-grained pruning for faster queries. - Micro-partitions can overlap in their range of values, which, combined with their uniformly small size, helps prevent skew. - Columns are stored independently within micro-partitions, often referred to as columnar storage. This enables efficient scanning of individual columns; only the columns referenced by a query are scanned. So - A is False (3rd point) - D is False because the sort is decided from "how the data is inserted" - E is False (1st point)

upvoted 8 times

**_yyukta** Most Recent 2 years, 2 months ago

BC are correct

upvoted 1 times

**Marge23** 2 years, 7 months ago

correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**Franzoz** 2 years, 10 months ago

Micro-partitions can overlap. No relationship to Time Travel. Sort is based on Data Ans = C D (Though rows are mapped not not sorted)

upvoted 1 times

**sakis213** 3 years, 4 months ago

B & C A.CAN OVERLAP D.micropartitions are created based on the order of the data loaded/inserted. Data are not sorted within micropart. E.are automatically derived based on the data, no schemas

upvoted 3 times

**KeshavaMugulur** 3 years, 4 months ago

Answer is A & B https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions.html

upvoted 1 times

**halol** 3 years, 4 months ago

C&D I think

upvoted 3 times

Question #295 *Topic 1*

Which data type can be used to store geospatial data in Snowflake?

A. Variant

B. Object

C. Geometry

D. Geography

**Suggested Answer:** *D*

*Community vote distribution*

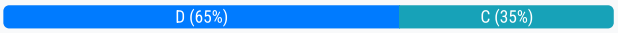


**Kate0204** Highly Voted 3 years, 5 months ago

Snowflake provides the following data types for geospatial data: The GEOGRAPHY data type, which models Earth as though it were a perfect sphere. The GEOMETRY data type, which represents features in a planar (Euclidean, Cartesian) coordinate system.

upvoted 8 times

**Andywu213** Highly Voted 3 years, 4 months ago

Geometry is in preview... assume this test refers to GA .. so D makes sense

upvoted 7 times

**PrashantGupta1616** Most Recent 1 year, 11 months ago

The GEOGRAPHY data type is specifically designed for storing geospatial data, allowing for spatial queries and analysis.

upvoted 3 times

**_yyukta** 2 years, 2 months ago

D is correct

upvoted 4 times

**Marge23** 2 years, 7 months ago

Data Types Snowflake provides the following data types for geospatial data: The GEOGRAPHY data type, which models Earth as though it were a perfect sphere. The GEOMETRY data type, which represents features in a planar (Euclidean, Cartesian) coordinate system. https://docs.snowflake.com/en/sql-reference/data-types-geospatial

upvoted 2 times

**singhks** 2 years, 8 months ago

C and D: https://docs.snowflake.com/en/sql-reference/data-types-geospatial Snowflake provides the following data types for geospatial data: The GEOGRAPHY data type, which models Earth as though it were a perfect sphere. The GEOMETRY data type, which represents features in a planar (Euclidean, Cartesian) coordinate system.

upvoted 2 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 2 times

**Hopefan** 2 years, 11 months ago

The GEOGRAPHY and GEOMETRY data types can store the following types of geospatial objects: WKT / WKB / EWKT / EWKB / GeoJSON geospatial objects: Point. MultiPoint.

upvoted 1 times

**learner2023** 3 years ago

C&D https://docs.snowflake.com/en/sql-reference/data-types-geospatial

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

For geospatial data storage like longitude/latitude Geography is the option https://docs.snowflake.com/en/sql-reference/data-types-geospatial#:~:text=If%20you%20have,use%20geospatial%20functionality.

upvoted 2 times

**KarBiswa** 3 years, 1 month ago

Both Geography and Geometry are geospatial data types https://docs.snowflake.com/en/sql-reference/data-types-geospatial two options must have given

upvoted 2 times

**SnowProCertDec22** 3 years, 4 months ago

https://docs.snowflake.com/en/sql-reference/data-types-geospatial.html#geometry-data-typeThe GEOGRAPHY data type, which models Earth as though it were a perfect sphere. The GEOMETRY data type, which represents features in a planar (Euclidean, Cartesian) coordinate system. [only PREVIEW as of Mar-2023] => D is correct

upvoted 1 times

**SnowProCertDec22** 3 years, 4 months ago

**Selected Answer: C**

https://docs.snowflake.com/en/sql-reference/data-types-geospatial.html#geometry-data-type

upvoted 1 times

**EmiB** 3 years, 2 months ago

The GEOGRAPHY data type, which models Earth as though it were a perfect sphere. The GEOMETRY data type, which represents features in a planar (Euclidean, Cartesian) coordinate system. [only PREVIEW as of Mar-2023] => D is correct

upvoted 1 times

Question #296 *Topic 1*

If all virtual warehouse resources are maximized while processing a query workload, what will happen to new queries that are submitted to the warehouse?

A. All queries will terminate when the resources are maximized.

B. The warehouse will scale out automatically

C. The warehouse will move to a suspended state.

D. New queries will be queued and executed when capacity is available.

**Suggested Answer:** *D*

*Community vote distribution*

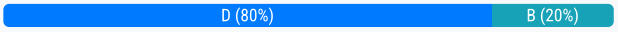


**_yyukta** Most Recent 2 years, 2 months ago

D is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**dmitriypo** 2 years, 12 months ago

D is correct

upvoted 1 times

**chizo** 3 years, 1 month ago

SO the keyword here maximized means MIN_CLUSTER_COUNT =MAX_CLUSTER_COUNT do the correct ans is D

upvoted 1 times

**shyemko** 3 years, 1 month ago

key word here is MAXIMIZED, means that we are here now in Multicluster warehouse with same value for parameters MIN_CLUSTER_COUNT and MAX_CLUSTER_COUNT

upvoted 1 times

**EmiB** 3 years, 1 month ago

D. New queries will be queued and executed when capacity is available.

upvoted 1 times

**Yogesh1984** 3 years, 2 months ago

D is correct

upvoted 3 times

Question #297 *Topic 1*

Masking policies can be applied to which of the following Snowflake objects? (Choose two.)

A. A materialized view

B. A stored procedure

C. A table

D. A stream

E. A pipe - E. A User-Defined Function (UDF)

**Suggested Answer:** *A, C*

*Community vote distribution*

AC (100%)

**halol** Highly Voted 3 years, 4 months ago

A&C I think

upvoted 13 times

**halol** Highly Voted 3 years, 4 months ago

A&C I thinkCreates a new masking policy in the current/specified schema or replaces an existing masking policy. After creating a masking policy, apply the masking policy to a column in a table using an ALTER TABLE … ALTER COLUMN command or a view using an ALTER VIEW command.
https://docs.snowflake.com/en/sql-reference/sql/create-masking-policy.html#normal-masking-policy

upvoted 13 times

**halol** 3 years, 4 months ago

Creates a new masking policy in the current/specified schema or replaces an existing masking policy. After creating a masking policy, apply the masking policy to a column in a table using an ALTER TABLE … ALTER COLUMN command or a view using an ALTER VIEW command.
https://docs.snowflake.com/en/sql-reference/sql/create-masking-policy.html#normal-masking-policy

upvoted 8 times

**_yyukta** Most Recent 2 years, 2 months ago

AC are correct

upvoted 1 times

**Marge23** 2 years, 7 months ago

applied to a table or view

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**DataEngineer7331** 3 years ago

A&C obviously

upvoted 1 times

**learner2023** 3 years ago

In Snowflake, masking policies are schema-level objects, which means a database and schema must exist in Snowflake before a masking policy can be applied to a column. Currently, Snowflake supports using Dynamic Data Masking on tables and views.

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

Table and Mat view can be applied with masking policies, https://docs.snowflake.com/en/user-guide/security-column-intro#label-security-column-intro-mv:~:text=Snowflake%20allows%20setting,materialized%20view%20columns.

upvoted 1 times

**EmiB** 3 years, 2 months ago

A&C Use the ALTER VIEW command to set/unset a masking policy, row access policy, or tag on/from a materialized view.

upvoted 1 times

**ppsingh001** 3 years, 3 months ago

A and C, for materialized views they are not allowed to be created on columns that has a masking policy. But you can apply a masking policy on materialized views.

upvoted 1 times

**BigDataBB** 3 years, 3 months ago

https://docs.snowflake.com/en/user-guide/security-column-intro.html#apply-a-conditional-masking-policy-on-a-column

upvoted 1 times

**yogen** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/security-column-intro.html#label-security-column-limitations Masking policy can be applied on Mat view columns as well

upvoted 1 times

**Kate0204** 3 years, 5 months ago

It's on schema, so table or stream, since they're within schema

upvoted 1 times

Question #298 *Topic 1*

What actions are supported by Snowflake resource monitors? (Choose two.)

A. Alert

B. Notify

C. Notify and suspend

D. Abort

E. Suspend immediately

**Suggested Answer:** *B, C*

*Community vote distribution*

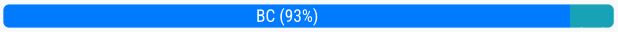


**halol** Highly Voted 3 years, 4 months ago

B&C https://docs.snowflake.com/en/user-guide/resource-monitors.html#actions Notify & Suspend Notify & Suspend Immediately Notify

upvoted 16 times

**KeshavaMugulur** Highly Voted 3 years, 4 months ago

B C & E is the answer

upvoted 5 times

**KeshavaMugulur** 3 years, 4 months ago

B C & E is the answerE. Suspend immediately >> sorry only "Notify & Suspend Immediately"

upvoted 5 times

**EmiB** 3 years, 1 month ago

E. Suspend immediately >> sorry only "Notify & Suspend Immediately"

upvoted 3 times

**PrasantSadangi** Most Recent 1 year, 7 months ago

Notify Notify & Suspend Notify & Suspend Immediately

upvoted 1 times

**_yyukta** 2 years, 2 months ago

B and C are correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**learner2023** 3 years ago

Resource monitors support the following actions: Notify & Suspend Send a notification (to all account administrators with notifications enabled) and suspend all assigned warehouses after all statements being executed by the warehouse(s) have completed. Notify & Suspend Immediately Send a notification (to all account administrators with notifications enabled) and suspend all assigned warehouses immediately, which cancels any statements being executed by the warehouses at the time. Notify Perform no action, but send an alert notification (to all account administrators with notifications enabled).

upvoted 4 times

**shyemko** 3 years, 1 month ago

There are 3 actions : - Notify - Notify & Suspend - Notify & Suspend immediatly

upvoted 2 times

**cgapperi** 3 years, 2 months ago

My read on this is B & C, because there is not a Suspend Immediately action. It is a Notify and Suspend Immediately action.

upvoted 2 times

**N89** 3 years, 2 months ago

CE represents exhaustive responsibilities

upvoted 1 times

**BigDataBB** 3 years, 3 months ago

B&C Notify Notify & Suspend Notify & Suspend Immediately https://docs.snowflake.com/en/user-guide/resource-monitors.html#actions

upvoted 2 times

**Apo1021** 3 years, 3 months ago

Should be B & C. notify notify & suspend notify & suspend immediately

upvoted 2 times

**paknadeem** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/resource-monitors.html#actions

upvoted 1 times

**paknadeem** 3 years, 4 months ago

The correct answer is B, E

upvoted 1 times

**Kate0204** 3 years, 5 months ago

Resource monitors support the following actions: Notify & Suspend Send a notification (to all account administrators with notifications enabled) and suspend all assigned warehouses after all statements being executed by the warehouse(s) have completed. Notify & Suspend Immediately Send a notification (to all account administrators with notifications enabled) and suspend all assigned warehouses immediately, which cancels any statements being executed by the warehouses at the time. Notify Perform no action, but send an alert notification (to all account administrators with notifications enabled).

upvoted 3 times

Question #299 *Topic 1*

A user executes the following SQL query:

create table SALES_BKP like SALES;

What are the cost implications for processing this query?

A. Processing costs will be generated based on how long the query takes.

B. Storage costs will be generated based on the size of the data.

C. No costs will be incurred as the query will use metadata.

D. The cost for running the virtual warehouse will be charged by the second.

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**halol** Highly Voted 3 years, 4 months ago

Snowflake SQL provides “CREATE TABLE LIKE” statement to create a new table with just the structure of the existing table without copying the data.
So I think it's C

upvoted 9 times

**KeshavaMugulur** Highly Voted 3 years, 4 months ago

C is right.

upvoted 9 times

**_yyukta** Most Recent 2 years, 2 months ago

C is correct

upvoted 1 times

**Marge23** 2 years, 7 months ago

CREATE TABLE … LIKE (creates an empty copy of an existing table) https://docs.snowflake.com/en/sql-reference/sql/create-table

upvoted 2 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**shyemko** 3 years, 1 month ago

CREATE TABLE … LIKE (creates an empty copy of an existing table) source : https://docs.snowflake.com/en/sql-reference/sql/create-table

upvoted 2 times

**tpiscaglia** 3 years, 1 month ago

CREATE TABLE LIKE creates an empty copy. It should use only metadata to create so it shouldn't need a warehouse/compute or storage.
https://docs.snowflake.com/en/sql-reference/sql/create-table#create-table-like

upvoted 1 times

**EmiB** 3 years, 2 months ago

CREATE TABLE … AS SELECT (creates a populated table; also referred to as CTAS) CREATE TABLE … USING TEMPLATE (creates a table with the column definitions derived from a set of staged files) CREATE TABLE … LIKE (creates an empty copy of an existing table) CREATE TABLE … CLONE (creates a clone of an existing table)

upvoted 4 times

**OTE** 3 years, 2 months ago

should be C

upvoted 1 times

Question #300 *Topic 1*

What is the maximum length of time travel available in the Snowflake Standard Edition?

A. 1 Day

B. 7 Days

C. 30 Days

D. 90 Days

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Kate0204** Highly Voted 3 years, 5 months ago

Only 1 day

upvoted 11 times

**_yyukta** Most Recent 2 years, 2 months ago

A. 1 Day

upvoted 1 times

**azlovemovie** 2 years, 4 months ago

1 day https://docs.snowflake.com/en/user-guide/data-time-travel#data-retention-period

upvoted 2 times

**Marge23** 2 years, 7 months ago

0-1 day only for Standard Edition

upvoted 1 times

**ntechkuma** 3 years, 1 month ago

1 day for std edition

upvoted 2 times

**mighty_kafta** 3 years, 2 months ago

A is correct

upvoted 4 times

**DK_1** 3 years, 2 months ago

only 1 day

upvoted 1 times

**Yogesh1984** 3 years, 2 months ago

A is correct only. 90 days in enterprise edition only.

upvoted 3 times

**fahfouhi94** 3 years, 3 months ago

one day only

upvoted 1 times

**Madhavan71** 3 years, 4 months ago

only 1 Day

upvoted 2 times

Question #301 *Topic 1*

What happens when an external or an internal stage is dropped? (Choose two.)

A. When dropping an external stage, the files are not removed and only the stage is dropped.

B. When dropping an external stage, both the stage and the files within the stage are removed.

C. When dropping an internal stage, the files are deleted with the stage and the files are recoverable.

D. When dropping an internal stage, the files are deleted with the stage and the files are not recoverable.

E. When dropping an internal stage, only selected files are deleted with the stage and are not recoverable.

**Suggested Answer:** *A, D*

*Community vote distribution*

AD (100%)

**SV1122** Highly Voted 3 years, 4 months ago

For an internal stage, all of the files in the stage are purged from Snowflake, regardless of their load status. This prevents the files from continuing to using storage and, consequently, accruing storage charges. However, this also means that the staged files cannot be recovered after a stage is dropped. For an external stage, only the stage itself is dropped; any data files in the referenced external location (Amazon S3, Google Cloud Storage, or Microsoft Azure) are not removed. https://docs.snowflake.com/en/sql-reference/sql/drop-stage.html#drop-stage

upvoted 9 times

**_yyukta** Most Recent 2 years, 2 months ago

AD are correct

upvoted 1 times

**Marge23** 2 years, 7 months ago

correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**imhm24** 2 years, 11 months ago

per the docs explain

upvoted 2 times

**sakis213** 3 years, 4 months ago

per the docs explain

upvoted 2 times

**halol** 3 years, 4 months ago

https://docs.snowflake.com/en/sql-reference/sql/drop-stage.html#drop-stage

upvoted 2 times

Question #302 *Topic 1*

A user has 10 files in a stage containing new customer data. The ingest operation completes with no errors, using the following command:

COPY INTO my_table FROM @my_stage;

The next day the user adds 10 files to the stage so that now the stage contains a mixture of new customer data and updates to the previous data. The user did not remove the 10 original files.

If the user runs the same COPY INTO command what will happen?

A. All data from all of the files on the stage will be appended to the table.

B. Only data about new customers from the new files will be appended to the table.

C. The operation will fail with the error UNCERTAIN FILES IN STAGE.

D. All data from only the newly-added files will be appended to the table.

**Suggested Answer:** *D*

*Community vote distribution*

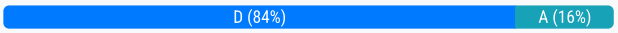


**MultiCloudIronMan** 1 year, 9 months ago

Selected Answer: A

According to the Snowflake documentation, when you run the COPY INTO command again without removing the original files, all data from all of the files on the stage will be appended to the table

upvoted 1 times

**Santa9os** Most Recent 1 year, 2 months ago

By default, COPY INTO tracks files that have been successfully loaded in the past. If you rerun the same command on the stage, Snowflake will NOT reload previously loaded files, unless you use the FORCE = TRUE option. Therefore, only the new files will be considered.

upvoted 1 times

**MultiCloudIronMan** 1 year, 9 months ago

According to the Snowflake documentation, when you run the COPY INTO command again without removing the original files, all data from all of the files on the stage will be appended to the tableI really dont understand the question

upvoted 1 times

**MultiCloudIronMan** 1 year, 8 months ago

I really dont understand the question

upvoted 2 times

**6c79d6f** 2 years ago

Once the files are copied into a table, they cannot be copied again in the next 64 days because of the files' metadata. If this option is true, it loads all files, regardless of whether they've been loaded previously and have not changed since they were loaded.
https://plazagonzalo.medium.com/snowflake-snowpro-data-loading-f88afaef7eb9#cf88

upvoted 1 times

**Heetec** 2 years, 7 months ago

what s the difference between B) and D) the meaning is the same, only data from new files will be added...

upvoted 1 times

**_yyukta** 2 years, 2 months ago

D is correct

upvoted 2 times

**Heetec** 2 years, 7 months ago

what s the difference between B) and D) the meaning is the same, only data from new files will be added...B: the new files could contains both old and new data.B Mentions only new customers from the new files meaning any duplicates in new files will be removed while loading which is not true. All the data from new files will be loaded as LOAD HISTORY is maintained by snowflake. So D is correct answer with this question.

upvoted 1 times

**gizzamo** 2 years, 4 months ago

B: the new files could contains both old and new data.

upvoted 1 times

**VASI250** 2 years, 4 months ago

B Mentions only new customers from the new files meaning any duplicates in new files will be removed while loading which is not true. All the data from new files will be loaded as LOAD HISTORY is maintained by snowflake. So D is correct answer with this question.

upvoted 1 times

**Marge23** 2 years, 7 months ago

D is correct Loading Older Files¶ This section describes how the COPY INTO command prevents data duplication differently based on whether the load status for a file is known or unknown. If you partition your data in stages using logical, granular paths by date (as recommended in Organizing Data by Path) and load data within a short period of time after staging it, this section largely does not apply to you. However, if the COPY command skips older files (i.e. historical data files) in a data load, this section describes how to bypass the default behavior. https://docs.snowflake.com/en/user-guide/data-load-considerations-load#executing-parallel-copy-statements-that-reference-the-same-data-files

upvoted 1 times

**EmiB** 3 years, 1 month ago

D - only the new files will be appended B - SF doesn't know which customer is new (COPY doesn't care about the meaning of data, care about files (file names) which are new (no kept in metadata as loaded)

upvoted 3 times

**fahfouhi94** 3 years, 3 months ago

D is correct answer

upvoted 2 times

**sakis213** 3 years, 4 months ago

only the new files will be appended

upvoted 2 times

**AS314** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/data-load-considerations-load.html#executing-parallel-copy-statements-that-reference-the-same-data-files

upvoted 4 times

**halol** 3 years, 4 months ago

Selected Answer: A

https://community.snowflake.com/s/question/0D50Z00009T3JWXSA3/upload-data-to-table-and-avoid-duplicated-data-rows

upvoted 2 times

**yogen** 3 years, 4 months ago

D - COPY command maintains historic load metadata with target table , so day 1 , 10 files will will not be loaded again

upvoted 3 times

**halol** 3 years, 4 months ago

https://community.snowflake.com/s/question/0D50Z00009T3JWXSA3/upload-data-to-table-and-avoid-duplicated-data-rowsCopy into not reload the same file more times if you don't set force = true https://docs.snowflake.com/en/sql-reference/sql/copy-into-table.html#reloading-files - Reloading Files Add FORCE = TRUE to a COPY command to reload (duplicate) data from a set of staged data files that have not changed (i.e. have the same checksum as when they were first loaded).

upvoted 2 times

**BigDataBB** 3 years, 3 months ago

Copy into not reload the same file more times if you don't set force = true https://docs.snowflake.com/en/sql-reference/sql/copy-into-table.html#reloading-files - Reloading Files Add FORCE = TRUE to a COPY command to reload (duplicate) data from a set of staged data files that have not changed (i.e. have the same checksum as when they were first loaded).

upvoted 4 times

Question #303 *Topic 1*

Which parameter can be used to instruct a COPY command to verify data files instead of loading them into a specified table?

A. STRIP_NULL_VALUES

B. SKIP_BYTE_ORDER_MARK

C. REPLACE_INVALID_CHARACTERS

D. VALIDATION_MODE

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**_yyukta** Most Recent 2 years, 2 months ago

D. VALIDATION_MODE

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**c0d3g** 2 years, 10 months ago

VALIDATION_MODE

upvoted 1 times

**SV1122** 3 years, 4 months ago

VALIDATION_MODE: This instructs the command to validate the data files instead of loading them into target tables and allows you to perform the dry run to ensure the fail-safe delivery of data.Dec

upvoted 3 times

Question #304 *Topic 1*

Which of the following SQL statements will list the version of the drivers currently being used?

A. Execute SELECT CURRENT_ODBC_CLIENT(); from the Web UI

B. Execute SELECT CURRENT_JDBC_VERSION(); from SnowSQL

C. Execute SELECT CURRENT_CLIENT(); from an application

D. Execute SELECT CURRENT_VERSION(); from the Python Connector

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**PrasantSadangi** Most Recent 1 year, 7 months ago

/*Returns the Current Client Version */ SELECT CURRENT_CLIENT(); /*Returns the Current Snowflake Version */ SELECT CURRENT_VERSION();

upvoted 1 times

**_yyukta** 2 years, 2 months ago

C is correct

upvoted 1 times

**0e504b5** 2 years, 3 months ago

https://docs.snowflake.com/en/sql-reference/functions/current_client

upvoted 1 times

**e_l_i_a_s** 2 years, 3 months ago

CURRENT_CLIENT Returns the version of the client from which the function was called. If called from an application using the JDBC or ODBC driver to connect to Snowflake, returns the version of the driver.

upvoted 1 times

**VASI250** 2 years, 4 months ago

The correct option for listing the version of the drivers currently being used in Snowflake is: B. Execute SELECT CURRENT_JDBC_VERSION(); from SnowSQL The CURRENT_JDBC_VERSION() function is used to retrieve the version of the JDBC driver in use. SnowSQL is a command-line client for Snowflake, and you can execute SQL statements directly using it. The other options do not directly provide information about the driver version: A. SELECT CURRENT_ODBC_CLIENT(); is used to get information about the ODBC client, not the version of the ODBC driver. C. SELECT CURRENT_CLIENT(); provides information about the client environment but not the driver version. D. SELECT CURRENT_VERSION(); is a generic function that provides information about the version of the Snowflake service itself, not the driver version. So, option B is the correct one for obtaining the version of the JDBC driver in use

upvoted 1 times

**444823b** 2 years, 4 months ago

SELECT CURRENT_CLIENT();

upvoted 1 times

**c0d3g** 2 years, 10 months ago

SELECT CURRENT_CLIENT();

upvoted 1 times

**dmitriypo** 2 years, 11 months ago

C is the correct answer. Just checked it.

upvoted 1 times

**EmiB** 3 years, 1 month ago

CURRENT_CLIENT() -- name and version of connected client CURRENT_VERSION() -- version of Snowflake

upvoted 4 times

**examed11** 3 years, 1 month ago

https://docs.snowflake.com/en/sql-reference/functions/current_client.html

upvoted 2 times

**AravindhTN** 3 years, 4 months ago

Answer C

upvoted 3 times

**harshagc180** 3 years, 5 months ago

Correct Answer C https://docs.snowflake.com/en/sql-reference/functions/current_client.html

upvoted 4 times

Question #305 *Topic 1*

Which Snowflake technique can be used to improve the performance of a query?

A. Clustering

B. Indexing

C. Fragmenting

D. Using INDEX_HINTS

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**harshagc180** Highly Voted 3 years, 5 months ago

Answer A , Clustering

upvoted 16 times

**tstander** Highly Voted 2 years, 10 months ago

This site would be useless without the community votes. *High5*

upvoted 14 times

**_yyukta** Most Recent 2 years, 2 months ago

A. Clustering

upvoted 1 times

**junaid2107** 2 years, 5 months ago

Indexing does not exist in Snowflake

upvoted 2 times

**Marge23** 2 years, 7 months ago

CORRECT

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

clustering

upvoted 1 times

**dw180** 2 years, 10 months ago

There is no index in snowflake

upvoted 1 times

**aroma** 3 years ago

Clustering

upvoted 1 times

**EmiB** 3 years, 1 month ago

A. Clustering

upvoted 1 times

**muks86** 3 years, 2 months ago

Answer is A, kindly update

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

A is correct an,swer

upvoted 1 times

**BigDataBB** 3 years, 3 months ago

A is correct

upvoted 1 times

**Apo1021** 3 years, 3 months ago

A is correct

upvoted 1 times

Question #306 *Topic 1*

What happens to the shared objects for users in a consumer account from a share, once a database has been created in that account?

A. The shared objects are transferred.

B. The shared objects are copied.

C. The shared objects become accessible.

D. The shared objects can be re-shared.

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**_yyukta** Most Recent 2 years, 2 months ago

C is correct

upvoted 1 times

**Marge23** 2 years, 7 months ago

CORRECT

upvoted 1 times

**singhks** 2 years, 8 months ago

C : Shared objects become accessible.

upvoted 2 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**aroma** 3 years ago

Database become accessible

upvoted 2 times

**fahfouhi94** 3 years, 3 months ago

Selected Answer: C

C : database become accessible

upvoted 3 times

**fahfouhi94** 3 years, 3 months ago

C : database become accessibleI mean db object

upvoted 3 times

**fahfouhi94** 3 years, 3 months ago

I mean db object

upvoted 2 times

**miadiesel** 3 years, 3 months ago

Once a database is created (in a consumer account) from a share, all the shared objects are accessible to users in the consumer account:
https://docs.snowflake.com/en/user-guide/data-sharing-intro.html#what-is-a-share

upvoted 4 times

**halol** 3 years, 4 months ago

C https://docs.snowflake.com/en/user-guide/data-sharing-intro.html#how-does-secure-data-sharing-work

upvoted 2 times

**KeshavaMugulur** 3 years, 4 months ago

Answer is C

upvoted 2 times

Question #307 *Topic 1*

Using variables in Snowflake is denoted by using which SQL character?

A. @

B. &

C. $

D. #

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**SV1122** Highly Voted 3 years, 4 months ago

set (min, max)=(40, 70); select $min; select avg(salary) from emp where age between $min and $max; https://docs.snowflake.com/en/sql-reference/session-variables.html#:~:text=Using%20Variables%20in%20SQL,-Variables%20can%20be&text=To%20distinguish%20them%20from%20bind,be%20prefixed%20with%20a%20%24%20sign.&text=Because%20the%20%24%20sign%20is%20the,character%20when%20used%20in%20identifiers.

upvoted 5 times

**_yyukta** Most Recent 2 years, 2 months ago

C is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**Madhavan71** 3 years, 4 months ago

Selected Answer: C

upvoted 1 times

**Rob__C** 3 years, 4 months ago

https://docs.snowflake.com/en/sql-reference/session-variables.html#variable-identifiers

upvoted 1 times

**halol** 3 years, 4 months ago

correct

upvoted 1 times

Question #308 *Topic 1*

Which commands should be used to grant the privilege allowing a role to select data from all current tables and any tables that will be created later in a schema? (Choose two.)

A. grant USAGE on all tables in schema DB1.SCHEMA to role MYROLE;

B. grant USAGE on future tables in schema DB1.SCHEMA to role MYROLE;

C. grant SELECT on all tables in schema DB1.SCHEMA to role MYROLE;

D. grant SELECT on future tables in schema DB1.SCHEMA to role MYROLE;

E. grant SELECT on all tables in database DB1 to role MYROLE;

F. grant SELECT on future tables in database DB1 to role MYROLE;

**Suggested Answer:** *C, D*

*Community vote distribution*

CD (100%)

**arpit_dataguy** Highly Voted 3 years, 4 months ago

Changing to C,D

upvoted 10 times

**examed11** Highly Voted 3 years, 1 month ago

Asking two separate things

upvoted 5 times

**_yyukta** Most Recent 2 years, 2 months ago

CD are correct

upvoted 1 times

**VenkataPhaniPavanKumar** 2 years, 3 months ago

Correct Answer: C & D GRANT SELECT ON ALL TABLES IN SCHEMA TO ROLE ; GRANT SELECT ON FUTURE TABLES IN SCHEMA TO ROLE ; Source : https://docs.snowflake.com/en/sql-reference/sql/grant-privilege

upvoted 1 times

**Afzy** 2 years, 6 months ago

They are asking for select only not modify so A will not be an answer, c & d os correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**Jtravismiller** 3 years, 3 months ago

A,C - You need Grant Usage before Grant Select. Grant Select will give you access to all current and future tables (without needing to specify future). Granting future is only when you need to restrict access to current tables.).

upvoted 2 times

**sailoo** 3 years, 3 months ago

A,D is the correct answer

upvoted 1 times

**sailoo** 3 years, 3 months ago

A,D is the correct answerbreak this down into 1. select data from all current tables (A) 2. any tables that will be created later in a schema (D)changing to C,D

upvoted 1 times

**sailoo** 3 years, 3 months ago

break this down into 1. select data from all current tables (A) 2. any tables that will be created later in a schema (D)changing to C,D

upvoted 1 times

**sailoo** 3 years, 3 months ago

break this down into 1. select data from all current tables (A) 2. any tables that will be created later in a schema (D)

upvoted 1 times

**sailoo** 3 years, 3 months ago

changing to C,D

upvoted 1 times

**Apo1021** 3 years, 3 months ago

I think C, D

upvoted 2 times

**arpit_dataguy** 3 years, 4 months ago

Answer--A,D

upvoted 1 times

**Rob__C** 3 years, 4 months ago

https://community.snowflake.com/s/article/How-to-grant-select-on-all-future-tables-in-a-schema-and-database-level

upvoted 2 times

Question #309 *Topic 1*

How can a user change which columns are referenced in a view?

A. Modify the columns in the underlying table

B. Use the ALTER VIEW command to update the view

C. Recreate the view with the required changes

D. Materialize the view to perform the changes

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**halol** Highly Voted 3 years, 4 months ago

Answer is C: https://docs.snowflake.com/en/sql-reference/sql/alter-view.html#alter-view

upvoted 11 times

**PrasantSadangi** Most Recent 1 year, 7 months ago

ALTER VIEW Modifies the properties for an existing view. Currently the only supported operations are: Renaming a view. Converting to (or reverting from) a secure view. Adding, overwriting, removing a comment for a view. Note that you cannot use this command to change the definition for a view. To change the view definition, you must drop the view and then recreate it.

upvoted 2 times

**_yyukta** 2 years, 2 months ago

C is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**c0d3g** 2 years, 10 months ago

To change the view definition, you must drop the view and then recreate it.

upvoted 1 times

**srj_mehta** 2 years, 10 months ago

C: ALTER VIEW Modifies the properties for an existing view. Currently the only supported operations are: Renaming a view. Converting to (or reverting from) a secure view. Adding, overwriting, removing a comment for a view. Note that you cannot use this command to change the definition for a view. To change the view definition, you must drop the view and then recreate it.

upvoted 2 times

**KarBiswa** 3 years, 1 month ago

only need to change the underlying SQL views are always replacable

upvoted 2 times

**examed11** 3 years, 1 month ago

Note that you cannot use this command to change the definition for a view. To change the view definition, you must drop the view and then recreate it.No I think we can create and replave the view, with modified changes no need to drop the view.

upvoted 1 times

**examed11** 3 years, 1 month ago

Selected Answer: C

Note that you cannot use this command to change the definition for a view. To change the view definition, you must drop the view and then recreate it.

upvoted 1 times

**akhandelwal680** 2 years, 11 months ago

No I think we can create and replave the view, with modified changes no need to drop the view.

upvoted 1 times

**SV1122** 3 years, 4 months ago

Currently the only supported operations are: Renaming a view. Converting to (or reverting from) a secure view. Adding, overwriting, removing a comment for a view. Note that you cannot use this command to change the definition for a view. To change the view definition, you must drop the view and then recreate it. https://docs.snowflake.com/en/sql-reference/sql/alter-view.html#alter-view

upvoted 3 times

Question #310 *Topic 1*

Which statement describes pruning?

A. The filtering or disregarding of micro-partitions that are not needed to return a query.

B. The return of micro-partitions values that overlap with each other to reduce a query's runtime.

C. A service that is handled by the Snowflake Cloud Services layer to optimize caching.

D. The ability to allow the result of a query to be accessed as if it were a table.

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**_yyukta** Most Recent 2 years, 2 months ago

A is correct

upvoted 1 times

**MultiCloudIronMan** 3 years ago

A is correct

upvoted 2 times

**sakis213** 3 years, 4 months ago

A is correct

upvoted 3 times

Question #311 *Topic 1*

Which SQL command can be used to see the CREATE definition of a masking policy?

A. SHOW MASKING POLICIES

B. DESCRIBE MASKING POLICY

C. GET_DDL

D. LIST MASKING POLICIES

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**AS314** Highly Voted 3 years, 3 months ago

This is the result of B https://docs.snowflake.com/en/sql-reference/sql/desc-masking-policy.html#example So B can't be the answer

And this is the result of C https://docs.snowflake.com/en/sql-reference/functions/get_ddl.html#examples So C is the answer

upvoted 11 times

**AS314** Highly Voted 3 years, 3 months ago

This is the result of B https://docs.snowflake.com/en/sql-reference/sql/desc-masking-policy.html#example So B can't be the answer And this is the result of C https://docs.snowflake.com/en/sql-reference/functions/get_ddl.html#examples So C is the answerB result is | 1 | SSN_MASK | (VAL VARCHAR) | VARCHAR(16777216) | case when current_role() in ('ANALYST') then val else '*********' end | C result is | create or replace view BOOKS_VIEW as select title, author from books_table; |

upvoted 11 times

**Def21** 2 years, 10 months ago

B result is | 1 | SSN_MASK | (VAL VARCHAR) | VARCHAR(16777216) | case when current_role() in ('ANALYST') then val else '*********' end | C result is | create or replace view BOOKS_VIEW as select title, author from books_table; |

upvoted 2 times

**PrashantGupta1616** Most Recent 1 year, 11 months ago

SELECT GET_DDL('MASKING POLICY', 'your_masking_policy_name');

upvoted 1 times

**_yyukta** 2 years, 2 months ago

C. GET_DDL

upvoted 1 times

**0e504b5** 2 years, 3 months ago

C https://docs.snowflake.com/en/sql-reference/functions/get_ddl Returns a DDL statement that can be used to recreate the specified object. GET_DDL currently supports the following object types: Policies (see CREATE MASKING POLICY , CREATE PASSWORD POLICY , CREATE ROW ACCESS POLICY , and CREATE SESSION POLICY , CREATE AUTHENTICATION POLICY)

upvoted 1 times

**Bujji1234** 2 years, 3 months ago

B is the answer if you want to replace an existing masking policy and need to see the current definition of the policy, call the GET_DDL function or run the DESCRIBE MASKING POLICY command.

upvoted 1 times

**Pradeep_7869** 2 years, 6 months ago

B is the answer

upvoted 1 times

**surya_lolla** 2 years, 9 months ago

https://docs.snowflake.com/en/sql-reference/sql/desc-masking-policy DESC MASKING POLICY

upvoted 1 times

**SatyamKishore** 2 years, 11 months ago

Going by snowflake documentation, seems both b & c is correct. https://docs.snowflake.com/en/sql-reference/sql/create-masking-policy#usage-notes

upvoted 2 times

**sailoo** 3 years, 3 months ago

B is the correct answer

upvoted 1 times

**sakis213** 3 years, 4 months ago

should be GET_DDL, describe returns only the body, not the whole create statement definition , for example input & return column data types values will be missing

upvoted 4 times

**yogen** 3 years, 4 months ago

C for Sure

upvoted 1 times

**scoob** 3 years, 4 months ago

C - GET_DDL returns the create statement to recreate the object. DESCRIBE will show the sql behind the policy but not in the form of a create statement.

upvoted 4 times

**halol** 3 years, 4 months ago

I think B&C https://docs.snowflake.com/en/sql-reference/sql/create-masking-policy.html#usage-notes

upvoted 3 times

**harshagc180** 3 years, 5 months ago

Correct answer DESCRIBE MASKING POLICY since the question is about command. If asked about the function then GET_DDL

upvoted 4 times

Question #312 *Topic 1*

Which of the following is the Snowflake Account_Usage.Metering_History view used for?

A. Gathering the hourly credit usage for an account

B. Compiling an account's average cloud services cost over the previous month

C. Summarizing the throughput of Snowpipe costs for an account

D. Calculating the funds left on an account's contract

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**SV1122** Highly Voted 3 years, 4 months ago

The METERING_HISTORY view in the ACCOUNT_USAGE schema can be used to return the hourly credit usage for an account within the last 365 days (1 year). https://docs.snowflake.com/en/sql-reference/account-usage/metering_history.html

upvoted 5 times

**_yyukta** Most Recent 2 years, 2 months ago

A. Gathering the hourly credit usage for an account

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

correct answer

upvoted 2 times

**halol** 3 years, 4 months ago

A is correct: https://docs.snowflake.com/en/sql-reference/account-usage/metering_history.html#metering-history-view

upvoted 1 times

Question #313 *Topic 1*

Query parsing and compilation occurs in which architecture layer of the Snowflake Cloud Data Platform?

A. Cloud services layer

B. Compute layer

C. Storage layer

D. Cloud agnostic layer

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**harshagc180** Highly Voted 3 years, 5 months ago

Answer : Cloud Service Layer

upvoted 13 times

**_yyukta** Most Recent 2 years, 2 months ago

A. Cloud services layer

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

A: Cloud services layer

upvoted 2 times

**EmiB** 3 years, 1 month ago

A. Cloud services layer

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

cloud service layer

upvoted 2 times

**Rob__C** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/intro-key-concepts.html#cloud-services

upvoted 1 times

Question #314 *Topic 1*

Which of the following Snowflake objects can be shared using a secure share? (Choose two.)

A. Materialized views

B. Sequences

C. Procedures

D. Tables

E. Secure User Defined Functions (UDFs)

**Suggested Answer:** *D, E*

*Community vote distribution*

DE (100%)

**SV1122** Highly Voted 3 years, 4 months ago

The following Snowflake database objects can be shared: Tables External tables Secure views Secure materialized views Secure UDFs
https://docs.snowflake.com/en/user-guide/data-sharing-intro.html#introduction-to-secure-data-sharing

upvoted 9 times

**harshagc180** Highly Voted 3 years, 5 months ago

Correct Answer D and E https://docs.snowflake.com/en/user-guide/data-sharing-intro.html#

upvoted 8 times

**_yyukta** Most Recent 2 years, 2 months ago

DE are correct

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

https://docs.snowflake.com/en/user-guide/data-sharing-intro.html

upvoted 1 times

**halol** 3 years, 4 months ago

D&E https://docs.snowflake.com/en/user-guide/data-sharing-intro.html#introduction-to-secure-data-sharing

upvoted 4 times

Question #315 *Topic 1*

What happens to the underlying table data when a CLUSTER BY clause is added to a Snowflake table?

A. Data is hashed by the cluster key to facilitate fast searches for common data values

B. Larger micro-partitions are created for common data values to reduce the number of partitions that must be scanned

C. Smaller micro-partitions are created for common data values to allow for more parallelism

D. Data may be colocated by the cluster key within the micro-partitions to improve pruning performance

**Suggested Answer:** *D*

*Community vote distribution*

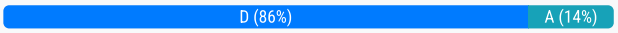


**halol** Highly Voted 3 years, 4 months ago

I think D: https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html#benefits-of-defining-clustering-keys-for-very-large-tables

upvoted 12 times

**SilviaS** Most Recent 10 months, 1 week ago

Not sure if D is the answer: "In Snowflake, defining a CLUSTER BY clause does not immediately reorganize the data. Instead, it enables the Automatic Clustering service, a background process that physically reorders the table data over time to group related rows into the same micro-partitions based on the specified clustering key.". There is a reorg hash over micropartitions. It's not that data MAY BE COLOCOATED, it IS colocated. thoughts?

upvoted 1 times

**_yyukta** 2 years, 2 months ago

D is correct

upvoted 1 times

**misaaa18** 2 years, 8 months ago

ans is D.

upvoted 1 times

**otsumy** 2 years, 8 months ago

must be D.

upvoted 1 times

**nagarajanbala** 3 years ago

D is correct

upvoted 2 times

**lordchao** 3 years, 1 month ago

D is correct

upvoted 2 times

**fahfouhi94** 3 years, 3 months ago

D is good answer

upvoted 2 times

**sakis213** 3 years, 4 months ago

that would be D

upvoted 2 times

**SV1122** 3 years, 4 months ago

A clustering key is a subset of columns in a table (or expressions on a table) that are explicitly designated to co-locate the data in the table in the same micro-partitions. This is useful for very large tables where the ordering was not ideal (at the time the data was inserted/loaded) or extensive DML has caused the table's natural clustering to degrade. https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html#benefits-of-defining-clustering-keys-for-very-large-tables

upvoted 4 times

Question #316 *Topic 1*

Which of the following conditions must be met in order to return results from the results cache? (Choose two.)

A. The user has the appropriate privileges on the objects associated with the query.

B. Micro-partitions have been reclustered since the query was last run.

C. The new query is run using the same virtual warehouse as the previous query.

D. The query includes a User Defined Function (UDF).

E. The query has been run within 24 hours of the previously-run query.

**Suggested Answer:** *A, E*

*Community vote distribution*

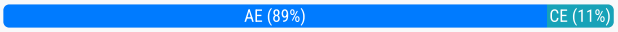


**halol** Highly Voted 3 years, 4 months ago

A&E as the below link: https://docs.snowflake.com/en/user-guide/querying-persisted-results.html#retrieval-optimization

upvoted 9 times

**halol** Highly Voted 3 years, 4 months ago

E is correct and maybe A If the same query is fired again in 24 hrs, it will not be COMPUTED, which means it will not be charged. and it is not affected by WH suspension. So as the previous question the same exact query will return the pre-computed results if the underlying data hasn't changed and the results were last accessed within previous 24-hour period

upvoted 6 times

**_yyukta** Most Recent 2 years, 2 months ago

AE are correct

upvoted 1 times

**Avi0407** 2 years, 6 months ago

C&E having access on an object might not give the same results... need to revisit the answer ..

upvoted 1 times

**examtopics_strata** 2 years, 9 months ago

A and E correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**c0d3g** 2 years, 10 months ago

The user has the appropriate privileges on the objects associated with the query. The query has been run within 24 hours of the previously-run query.

upvoted 1 times

**Hopefan** 3 years, 1 month ago

how about C and E ?

upvoted 1 times

**Hopefan** 3 years, 1 month ago

how about C and E ?it doesn't matter with the warehouse since the result cache returning is through cloud service layer.

upvoted 1 times

**cctv** 2 years, 8 months ago

it doesn't matter with the warehouse since the result cache returning is through cloud service layer.

upvoted 3 times

**fahfouhi94** 3 years, 3 months ago

A&E are the conditions

upvoted 3 times

**AlexbDku** 3 years, 4 months ago

Answer -- A&E

upvoted 1 times

**arpit_dataguy** 3 years, 4 months ago

Answer -- A&E

upvoted 3 times

Question #317 *Topic 1*

Which statement about billing applies to Snowflake credits?

A. Credits are billed per-minute with a 60-minute minimum.

B. Credits are used to pay for cloud data storage usage.

C. Credits are consumed based on the number of credits billed for each hour that a warehouse runs.

D. Credits are consumed based on the warehouse size and the time the warehouse is running.

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**arpit_dataguy** Highly Voted 3 years, 4 months ago

D is the answer. A virtual warehouse is one or more clusters of compute resources that enable executing queries, loading data, and performing other DML operations. Snowflake credits are used to pay for the processing time used by each virtual warehouse. Snowflake credits are charged based on the number of virtual warehouses you use, how long they run, and their size.

upvoted 12 times

**PrashantGupta1616** Most Recent 1 year, 11 months ago

Snowflake credits are indeed consumed based on the size of the virtual warehouse and the duration for which it is running. Billing is done on a per-second basis with a minimum of 60 seconds for each usage period.

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

Correct answer D

upvoted 2 times

**EmiB** 3 years, 1 month ago

D - based on the "warehouse size" and the "time the warehouse is running".

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

change to D

upvoted 1 times

**Hopefan** 3 years, 1 month ago

D is correct

upvoted 1 times

**paknadeem** 3 years, 4 months ago

I think the answer is D

upvoted 3 times

**halol** 3 years, 4 months ago

I think D https://docs.snowflake.com/en/user-guide/cost-understanding-compute.html#virtual-warehouse-credit-usage

upvoted 3 times

Question #318 *Topic 1*

A user needs to create a materialized view in the schema MYDB.MYSCHEMA.

Which statements will provide this access?

A. GRANT ROLE MYROLE TO USER USER1; CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO ROLE MYROLE;

B. GRANT ROLE MYROLE TO USER USER1; CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER USER1;

C. GRANT ROLE MYROLE TO USER USER1; CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER1;

D. GRANT ROLE MYROLE TO USER USER1; CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE;

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**halol** Highly Voted 3 years, 4 months ago

A is the correct answer https://docs.snowflake.com/en/sql-reference/sql/grant-privilege.html#examples

upvoted 10 times

**BigDataBB** Highly Voted 3 years, 4 months ago

in all commands it is missing the keyword GRANT before "CREATE MATERIALIZED VIEW" The answer A is correct but the answer D is correct too because the keyword ROLE is optional: https://docs.snowflake.com/en/sql-reference/sql/grant-privilege.html TO [ ROLE ] [ WITH GRANT OPTION ] So the question must allow for multiple answers.I forgot to put at the beginning of the sentence: "Premising that"

upvoted 6 times

**BigDataBB** 3 years, 4 months ago

in all commands it is missing the keyword GRANT before "CREATE MATERIALIZED VIEW" The answer A is correct but the answer D is correct too because the keyword ROLE is optional:

https://docs.snowflake.com/en/sql-reference/sql/grant-privilege.html TO [ ROLE ] [ WITH GRANT OPTION ]

So the question must allow for multiple answers.

upvoted 6 times

**BigDataBB** 3 years, 4 months ago

I forgot to put at the beginning of the sentence: "Premising that"

upvoted 1 times

**_yyukta** Most Recent 2 years, 2 months ago

A is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

A and D are correct

upvoted 2 times

**KnightVictor** 2 years, 11 months ago

A is correct

upvoted 2 times

**EmiB** 3 years, 1 month ago

A & D - both correct B & C - fail (only role can be granted to a user) GRANT CREATE MATERIALIZED VIEW ON SCHEMA mydb.myschema TO ROLE myrole; GRANT CREATE MATERIALIZED VIEW ON SCHEMA mydb.myschema TO USER user1; GRANT CREATE MATERIALIZED VIEW ON SCHEMA mydb.myschema TO user1; GRANT CREATE MATERIALIZED VIEW ON SCHEMA mydb.myschema TO myrole;

upvoted 4 times

**ibotz** 3 years, 4 months ago

grant create materialized view on schema to role ;

upvoted 2 times

Question #319 *Topic 1*

What is the purpose of multi-cluster virtual warehouses?

A. To create separate data warehouses to increase query optimization

B. To allow users the ability to choose the type of compute nodes that make up a virtual warehouse cluster

C. To eliminate or reduce queuing of concurrent queries

D. To allow the warehouse to resize automatically

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**fahfouhi94** Highly Voted 3 years, 3 months ago

answer is correct

upvoted 6 times

**_yyukta** Most Recent 2 years, 2 months ago

C. To eliminate or reduce queuing of concurrent queries

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

Question #320 *Topic 1*

Which of the following is a valid source for an external stage when the Snowflake account is located on Microsoft Azure?

A. An FTP server with TLS encryption

B. An HTTPS server with WebDAV

C. A Google Cloud storage bucket

D. A Windows server file share on Azure

**Suggested Answer:** *C*

*Community vote distribution*

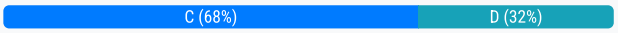


**ShagunMittal** Highly Voted 3 years, 1 month ago

Correct answer is C. Snowflake supports loading data from files staged in any of the following locations, regardless of the cloud platform for your Snowflake account: • Internal (i.e. Snowflake) stages • Amazon S3 • Google Cloud Storage • Microsoft Azure blob storage

upvoted 6 times

**devu70** Most Recent 1 year, 1 month ago

Snowflake is designed to be multi-cloud compatible, which means: You can create an external stage that points to a GCS bucket, regardless of whether your Snowflake account is on Azure, AWS, or GCP. This is done using a storage integration that securely connects Snowflake to the GCS bucket without exposing credentials.

upvoted 1 times

**ravinder_poonia1** 1 year, 8 months ago

In Snowflake, when the account is located on Microsoft Azure, a valid source for an external stage can be an Azure container or a folder path within an Azure container. This includes Azure Blob storage which is accessible via the azure:// endpoint. A Windows server file share on Azure, if configured properly, can be a valid source for staging data files for Snowflake. Options A, B, and C are not supported as direct sources for an external stage in Snowflake on Azure12.

upvoted 1 times

**PrashantGupta1616** 1 year, 11 months ago

Google storage is correct as AZURE file share is not supported with Snowflake staging.

upvoted 1 times

**_yyukta** 2 years, 2 months ago

C is correct

upvoted 2 times

**0e504b5** 2 years, 3 months ago

https://docs.snowflake.com/en/sql-reference/sql/create-stage External stage References data files stored in a location outside of Snowflake. Currently, the following cloud storage services are supported: Amazon S3 buckets Google Cloud Storage buckets Microsoft Azure containers

upvoted 1 times

**romega2** 2 years, 7 months ago

correct

upvoted 1 times

**Rajivnb** 2 years, 9 months ago

To my earlier question. The reason its Google storage being correct is, AZURE file share is not supported with Snowflake staging. Snowflake currently supports loading from blob storage only. Snowflake supports the following types of storage accounts. It does not contain the windows share in the list. Blob storage Data Lake Storage Gen2 General-purpose v1 General-purpose v2

upvoted 2 times

**Rajivnb** 2 years, 9 months ago

I dont understand why some of them say C. Though Snowflake sits on the Azure from the questions, the external stage is anything out of snowflake. It could be Azure File share or it could be Google bucket. can someone clarify?

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

MS AzureIts not D. Snowflake does not support Auzre Windows Share. It only supports BLOB

upvoted 2 times

**KnightVictor** 2 years, 11 months ago

Selected Answer: D

MS Azure

upvoted 2 times

**Rajivnb** 2 years, 6 months ago

Its not D. Snowflake does not support Auzre Windows Share. It only supports BLOB

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

changing the voting to C this is correct

upvoted 2 times

**KarBiswa** 3 years, 1 month ago

as it is MS based environment

upvoted 3 times

**EmiB** 3 years, 1 month ago

You can hosted SF at one provide but access data via ext. stage from all three (S3 Bucket, Azure Blob, GCP as well).

upvoted 1 times

**cgapperi** 3 years, 2 months ago

It cannot be C. All resources must be in the same cloud environment.https://docs.snowflake.com/en/user-guide/data-load-overview#external-stages Loading data from any of the following cloud storage services is supported regardless of the cloud platform that hosts your Snowflake account: Amazon S3 Google Cloud Storage Microsoft Azure

upvoted 3 times

**cgapperi** 3 years, 2 months ago

Selected Answer: D

It cannot be C. All resources must be in the same cloud environment.

upvoted 3 times

**okurok** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/data-load-overview#external-stages Loading data from any of the following cloud storage services is supported regardless of the cloud platform that hosts your Snowflake account: Amazon S3 Google Cloud Storage Microsoft Azure

upvoted 3 times

**ainouche928** 3 years, 2 months ago

The correct answer is D

upvoted 2 times

**kibic** 3 years, 3 months ago

C is correct

upvoted 3 times

Question #321 *Topic 1*

Which database objects can be shared with the Snowflake secure data sharing feature? (Choose two.)

A. Files

B. External tables

C. Secure User-Defined Functions (UDFs)

D. Sequences

E. Streams

**Suggested Answer:** *B, C*

*Community vote distribution*

BC (100%)

**_yyukta** Most Recent 2 years, 2 months ago

BC are correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**okurok** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/data-sharing-intro Secure Data Sharing lets you share selected objects in a database in your account with other Snowflake accounts. You can share the following Snowflake database objects: Tables External tables Secure views Secure materialized views Secure UDFs

upvoted 2 times

**Hopefan** 3 years, 1 month ago

BC correct

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

correct answer

upvoted 1 times

**SV1122** 3 years, 4 months ago

B. External tables C. Secure User-Defined Functions (UDFs)

upvoted 4 times

Question #322 *Topic 1*

Which statements reflect key functionalities of a Snowflake Data Exchange? (Choose two.)

A. If an account is enrolled with a Data Exchange, it will lose its access to the Snowflake Marketplace.

B. A Data Exchange allows groups of accounts to share data privately among the accounts.

C. A Data Exchange allows accounts to share data with third, non-Snowflake parties.

D. Data Exchange functionality is available by default in accounts using the Enterprise edition or higher.

E. The sharing of data in a Data Exchange is bidirectional. An account can be a provider for some datasets and a consumer for others.

**Suggested Answer:** *E*

*Community vote distribution*




**SV1122** Highly Voted 3 years, 4 months ago

B and E

upvoted 21 times

**Simbajagnole** Highly Voted 4 months ago

Selected Answer: A

A is right. Don't overthink it

upvoted 9 times

**harshagc180** 3 years, 4 months ago

Answer : B and C https://docs.snowflake.com/en/user-guide/data-exchange-benefits.htmlchanging answer to B & E

upvoted 6 times

**Matsalearnings** 4 months ago

seems i am already overthinking - why A? if another user is from diff account then how the results can be retrieved from result cache ?, Thanks

upvoted 2 times

**harshagc180** 3 years, 3 months ago

changing answer to B & E

upvoted 9 times

**18527ec** 4 months ago

Selected Answer: C

It another user executes the same query , it needs virtual warehouse

upvoted 1 times

**SergesA** 1 year, 9 months ago

Correct anwser B and E

upvoted 1 times

**_yyukta** 4 months ago

Selected Answer: A

A. No virtual warehouse will be used, data will be read from the result cache.

upvoted 1 times

**PrashantGupta1616** 1 year, 11 months ago

B and E

upvoted 1 times

**MultiCloudIronMan** Most Recent 4 months ago

Selected Answer: A

Correct

upvoted 1 times

**_yyukta** 2 years, 2 months ago

B. A Data Exchange allows groups of accounts to share data privately among the accounts.

upvoted 1 times

**majew** 2 years, 3 months ago

B and E

upvoted 1 times

**sandy91** 2 years, 3 months ago

Both B and E

upvoted 1 times

**gizzamo** 2 years, 4 months ago

B and E are correct. Question is asking us to choose 2!

upvoted 1 times

**Avi0407** 2 years, 6 months ago

E also one of the answer with B

upvoted 1 times

**Avi0407** 2 years, 6 months ago

B&E this fellow didn't give access to select 2

upvoted 1 times

**SCube** 2 years, 6 months ago

B & E as per teh definition of Data Exchange.

upvoted 1 times

**_Vicky** 2 years, 8 months ago

Both B and E are valid

upvoted 1 times

**conrad2023** 2 years, 8 months ago

B and E should be correct, but here you can only select one in the comment survey.

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

B and E are correct

upvoted 1 times

**BigDataBB** 3 years, 3 months ago

B&E https://docs.snowflake.com/en/user-guide/data-exchange-benefits.html

upvoted 4 times

Question #323 *Topic 1*

A Snowflake user executed a query and received the results. Another user executed the same query 4 hours later. The data had not changed.

What will occur?

A. No virtual warehouse will be used, data will be read from the result cache.

B. No virtual warehouse will be used, data will be read from the local disk cache.

C. The default virtual warehouse will be used to read all data.

D. The virtual warehouse that is defined at the session level will be used to read all data.

**Suggested Answer:** *A*

*Community vote distribution*

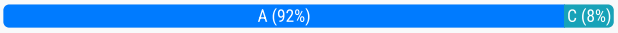


**Simbajagnole** Highly Voted 3 years, 2 months ago

A is right. Don't overthink itseems i am already overthinking - why A? if another user is from diff account then how the results can be retrieved from result cache ?, Thanks

upvoted 9 times

**Simbajagnole** Highly Voted 3 years, 2 months ago

Selected Answer: A

A is right. Don't overthink it

upvoted 9 times

**Matsalearnings** 2 years, 7 months ago

seems i am already overthinking - why A? if another user is from diff account then how the results can be retrieved from result cache ?, Thanks

upvoted 2 times

**18527ec** Most Recent 1 year, 5 months ago

It another user executes the same query , it needs virtual warehouse

upvoted 1 times

**_yyukta** 2 years, 2 months ago

A. No virtual warehouse will be used, data will be read from the result cache.

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

For sure

upvoted 1 times

Question #324 *Topic 1*

Which feature allows a user the ability to control the organization of data in a micro-partition?

A. Range Partitioning

B. Search Optimization Service

C. Automatic Clustering

D. Horizontal Partitioning

**Suggested Answer:** *C*

*Community vote distribution*

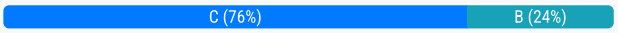


**halol** Highly Voted 3 years, 4 months ago

I think C

upvoted 15 times

**halol** Highly Voted 4 months ago

Selected Answer: D

https://docs.snowflake.com/en/user-guide/data-sharing-mutiple-db.html

upvoted 11 times

**Lematthew31** 2 years ago

IAutomatic Clustering allows users to define clustering keys on tables, which determines how data is physically organized within micro-partitions on the disk. This organization improves query performance by grouping similar data together, reducing the amount of data that needs to be scanned to satisfy a query

upvoted 1 times

**_yyukta** 4 months ago

Selected Answer: D

D is correct

upvoted 1 times

**_yyukta** 2 years, 2 months ago

C is correct

upvoted 1 times

**MultiCloudIronMan** 4 months ago

Selected Answer: D

correct

upvoted 1 times

**sandy91** 2 years, 3 months ago

C will be the right answer

upvoted 2 times

**KarBiswa** 4 months ago

Selected Answer: D

https://docs.snowflake.com/en/user-guide/data-sharing-mutiple-db

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

I think C is the answer, B is possible but C in this case

upvoted 2 times

**BigDataBB** Most Recent 4 months ago

Selected Answer: D

In addition to performing all the standard steps to share data, you must grant the REFERENCE_USAGE privilege on each database referenced by a secure view you wish to share. However, you do not need to grant REFERENCE_USAGE on the database where the secure view is created.

upvoted 1 times

**Def21** 2 years, 10 months ago

Automatic Clustering = Defining Clustering key "All you need to do is define a clustering key for each table" https://docs.snowflake.com/en/user-guide/tables-auto-reclustering

upvoted 3 times

**dmitriypo** 2 years, 11 months ago

C is correct

upvoted 1 times

**chizo** 3 years, 1 month ago

B is the correct ans

upvoted 1 times

**ShagunMittal** 3 years, 1 month ago

The search optimization service relies on a persistent data structure that serves as an optimized search access path. This search access path and the maintenance service are transparent to the user. You don't need to create a warehouse for the service that maintains the search access path. B cannot be the answer

upvoted 3 times

**nencheppan** 3 years, 1 month ago

As User allows its Search Optimization Service

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/tables-auto-reclustering

upvoted 3 times

**ShagunMittal** 3 years, 1 month ago

Search Optimization Service has nothing to do with the organization of data in micro-partitions. Its clustering only which organizes data in micro-partitionsBy users ability they mean creating a clustering key.

upvoted 4 times

**ShagunMittal** 3 years, 1 month ago

By users ability they mean creating a clustering key.

upvoted 1 times

**EmiB** 3 years, 1 month ago

B is correct as "allows user the ability to control the organization of data in a micro-partition" C - incorrect as "Manual Reclustering — Deprecated"

upvoted 1 times

**mayank** 3 years, 2 months ago

I think it is B SEARCH OPTIMIZATION SERVICE since the questions says "allows user the ability to control the organization of data in a micro-partition"

upvoted 3 times

**OTE** 3 years, 2 months ago

Should be C

upvoted 1 times

**fahfouhi94** 3 years, 2 months ago

https://docs.snowflake.com/en/user-guide/tables-auto-reclustering

upvoted 1 times

**BigDataBB** 3 years, 3 months ago

I think is B, because You can ask ti SF to "modify" the clustering of a table with the "clustering key" and the "Search Optimization Service" https://docs.snowflake.com/en/user-guide/search-optimization-service.html#managing-the-costs-of-the-search-optimization-service

upvoted 2 times

Question #325 *Topic 1*

Which privilege must be granted to a share to allow secure views the ability to reference data in multiple databases?

A. CREATE_SHARE on the account

B. SHARE on databases and schemas

C. SELECT on tables used by the secure view

D. REFERENCE_USAGE on databases

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**halol** Highly Voted 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/data-sharing-mutiple-db.html

upvoted 11 times

**_yyukta** Most Recent 2 years, 2 months ago

D is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/data-sharing-mutiple-db

upvoted 1 times

**BigDataBB** 3 years, 3 months ago

In addition to performing all the standard steps to share data, you must grant the REFERENCE_USAGE privilege on each database referenced by a secure view you wish to share. However, you do not need to grant REFERENCE_USAGE on the database where the secure view is created.

upvoted 1 times

Question #326 *Topic 1*

In which use case does Snowflake apply egress charges?

A. Data sharing within a specific region

B. Query result retrieval

C. Database replication

D. Loading data into Snowflake

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**harshagc180** Highly Voted 3 years, 4 months ago

Answer is Database Replication

upvoted 9 times

**BigDataBB** Highly Voted 3 years, 3 months ago

Snowflake charges a per-byte fee for data egress when users transfer data from a Snowflake account into a different region on the same cloud platform or into a completely different cloud platform. Data transfers within the same region are free. https://docs.snowflake.com/en/user-guide/cost-understanding-data-transfer.html

upvoted 6 times

**_yyukta** Most Recent 2 years, 2 months ago

C. Database replication

upvoted 1 times

**Afzy** 2 years, 6 months ago

Answer shod be A, egress charges means when data is transfered to another region or anothe cloud platform

upvoted 1 times

**otsumy** 2 years, 8 months ago

The answer must be A for sure

Contact your cloud storage provider (Amazon S3, Google Cloud Storage, or Microsoft Azure) to determine whether they apply data egress charges to transfer data from their network and region of origin to the cloud provider's network and region where your Snowflake account is hosted. https://docs.snowflake.com/en/user-guide/cost-understanding-data-transfer

upvoted 1 times

**otsumy** 2 years, 8 months ago

The answer must be A for sure Contact your cloud storage provider (Amazon S3, Google Cloud Storage, or Microsoft Azure) to determine whether they apply data egress charges to transfer data from their network and region of origin to the cloud provider's network and region where your Snowflake account is hosted. https://docs.snowflake.com/en/user-guide/cost-understanding-data-transferCan not be A Data sharing is not data transferring. With Secure Data Sharing, no actual data is copied or transferred between accounts. All sharing uses Snowflake's services layer and metadata store. Shared data does not take up any storage in a consumer account and therefore does not contribute to the consumer's monthly data storage charges. The only charges to consumers are for the compute resources (i.e. virtual warehouses) used to query the shared data. https://docs.snowflake.com/en/user-guide/data-sharing-introbut it share to different cloud platform or region, it might involve database replication, then it could be A

upvoted 1 times

**KindleAce** 2 years, 8 months ago

Can not be A Data sharing is not data transferring. With Secure Data Sharing, no actual data is copied or transferred between accounts. All sharing uses Snowflake's services layer and metadata store. Shared data does not take up any storage in a consumer account and therefore does not contribute to the consumer's monthly data storage charges. The only charges to consumers are for the compute resources (i.e. virtual warehouses) used to query the shared data. https://docs.snowflake.com/en/user-guide/data-sharing-intro

upvoted 1 times

**KindleAce** 2 years, 8 months ago

Can not be A Data sharing is not data transferring. With Secure Data Sharing, no actual data is copied or transferred between accounts. All sharing uses Snowflake's services layer and metadata store. Shared data does not take up any storage in a consumer account and therefore does not contribute to the consumer's monthly data storage charges. The only charges to consumers are for the compute resources (i.e. virtual warehouses) used to query the shared data. https://docs.snowflake.com/en/user-guide/data-sharing-introbut it share to different cloud platform or region, it might involve database replication, then it could be A

upvoted 1 times

**KindleAce** 2 years, 8 months ago

but it share to different cloud platform or region, it might involve database replication, then it could be A

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

For sure Database Replication

upvoted 1 times

**BigDataBB** 3 years, 3 months ago

https://docs.snowflake.com/en/user-guide/cost-understanding-data-transfer.html

upvoted 2 times

**Rob__C** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/cost-understanding-data-transfer.html

upvoted 2 times

Question #327 *Topic 1*

Which of the following compute resources or features are managed by Snowflake? (Choose two.)

A. Execute a COPY command

B. Updating data

C. Snowpipe

D. AUTOMATIC_CLUSTERING

E. Scaling up a warehouse

**Suggested Answer:** *C, D*

*Community vote distribution*




**SilviaS** Most Recent 10 months, 1 week ago

The question actually is which of these compute objects are serverless.

upvoted 1 times

**_yyukta** 2 years, 2 months ago

CD are correct

upvoted 2 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**learner2023** 3 years ago

https://docs.snowflake.com/en/user-guide/cost-understanding-compute

upvoted 2 times

**examed11** 3 years, 1 month ago

Serverless compute — Compute resources managed by Snowflake, rather than traditional virtual warehouses.

upvoted 2 times

**EmiB** 3 years, 1 month ago

C. Snowpipe E. Scaling up a warehouse D. AUTOMATIC_CLUSTERING - is not a feature, Automatic clustering is!

upvoted 2 times

**kibic** 3 years, 3 months ago

https://docs.snowflake.com/en/user-guide/cost-understanding-compute.html#serverless-features

upvoted 4 times

**BigDataBB** 3 years, 3 months ago

- Snowpipe uses Snowflake-supplied compute resources. - Automatic Clustering is the Snowflake service that seamlessly and continually manages all reclustering, as needed, of clustered tables.

upvoted 2 times

**paknadeem** 3 years, 4 months ago

i think the answer is D, E

upvoted 2 times

Question #328 *Topic 1*

A materialized view should be created when which of the following occurs? (Choose two.)

A. There is minimal cost associated with running the query.

B. The query consumes many compute resources every time it runs.

C. The base table gets updated frequently.

D. The query is highly optimized and does not consume many compute resources.

E. The results of the query do not change often and are used frequently.

**Suggested Answer:** *B, E*

*Community vote distribution*

BE (100%)

**halol** Highly Voted 3 years, 4 months ago

B&E https://docs.snowflake.com/en/user-guide/views-materialized.html#deciding-when-to-create-a-materialized-view

upvoted 12 times

**_yyukta** Most Recent 2 years, 2 months ago

BE are correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**Harpreetwal** 3 years ago

B&E Correct

upvoted 2 times

**cgapperi** 3 years, 2 months ago

"The query consumes a lot of resources. Typically, this means that the query consumes a lot of processing time or credits, but it could also mean that the query consumes a lot of storage space for intermediate results." This is why it is not D

upvoted 3 times

**yogendrat52** 3 years, 2 months ago

B and D is the right answer Create a materialized view when all of the following are true: The query results from the view don't change often. This almost always means that the underlying/base table for the view doesn't change often, or at least that the subset of base table rows used in the materialized view don't change often. The results of the view are used often (typically significantly more often than the query results change). The query consumes a lot of resources. Typically, this means that the query consumes a lot of processing time or credits, but it could also mean that the query consumes a lot of storage space for intermediate results.

upvoted 2 times

**kokosz** 3 years, 4 months ago

BE correct

upvoted 1 times

Question #329 *Topic 1*

What privilege should a user be granted to change permissions for new objects in a managed access schema?

A. Grant the OWNERSHIP privilege on the schema.

B. Grant the OWNERSHIP privilege on the database.

C. Grant the MANAGE GRANTS global privilege.

D. Grant ALL privileges on the schema.

**Suggested Answer:** *A*

*Community vote distribution*

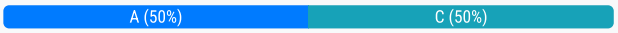


**AndroJS** Highly Voted 3 years, 8 months ago

It should be A because as both a & c answer are correct, the 'minimum' impacting option is Ownership

upvoted 5 times

**18527ec** Most Recent 1 year, 5 months ago

only schema owner can grant schema permissions

upvoted 1 times

**bor4un** 1 year, 9 months ago

https://docs.snowflake.com/en/sql-reference/sql/grant-ownership#usage-notes A role that has the MANAGE GRANTS privilege can transfer ownership of an object to any role; in contrast, a role that does not have the MANAGE GRANTS privilege can only transfer ownership from itself to a child role within the role hierarchy. The transfer of ownership (GRANT OWNERSHIP) only affects existing objects at the time the command is issued. Any objects created after the command is issued are owned by the role in use when the object is created.Hi I cannot find where you identified above explanation. Frankly, I never saw this idea related to future objects from the schema. I thought that owners ship of the schema will transform it in a regular schema (with DAC approach).

upvoted 1 times

**bor4un** 1 year, 9 months ago

Selected Answer: C

https://docs.snowflake.com/en/sql-reference/sql/grant-ownership#usage-notes

A role that has the MANAGE GRANTS privilege can transfer ownership of an object to any role; in contrast, a role that does not have the MANAGE GRANTS privilege can only transfer ownership from itself to a child role within the role hierarchy.

The transfer of ownership (GRANT OWNERSHIP) only affects existing objects at the time the command is issued. Any objects created after the command is issued are owned by the role in use when the object is created.

upvoted 1 times

**37ceea2** 1 year, 8 months ago

Hi I cannot find where you identified above explanation. Frankly, I never saw this idea related to future objects from the schema. I thought that owners ship of the schema will transform it in a regular schema (with DAC approach).

upvoted 1 times

**Qkel** 1 year, 11 months ago

A managed access schema is a way to centralize the management of access permissions for objects in a schema. This is done by limiting the ability to grant privileges to only the schema owner or roles with the MANAGE GRANTS privilege.

upvoted 1 times

**d22770a** 2 years ago

C is correct

upvoted 1 times

**vc5** 2 years, 7 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-configure

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A is correct

upvoted 2 times

**0e504b5** 2 years, 8 months ago

AC are both correct With managed access schemas, object owners lose the ability to make grant decisions. Only the schema owner (i.e. the role with the OWNERSHIP privilege on the schema) or a role with the MANAGE GRANTS privilege can grant privileges on objects in the schema, including future grants, centralizing privilege management.

upvoted 1 times

**0e504b5** 2 years, 9 months ago

https://docs.snowflake.com/en/sql-reference/sql/create-schema CREATE SCHEMA WITH MANAGED ACCESS Specifies a managed schema. Managed access schemas centralize privilege management with the schema owner. In regular schemas, the owner of an object (i.e. the role that has the OWNERSHIP privilege on the object) can grant further privileges on their objects to other roles. In managed schemas, the schema owner manages all privilege grants, including future grants, on objects in the schema. Object owners retain the OWNERSHIP privileges on the objects; however, only the schema owner can manage privilege grants on the objects.

upvoted 1 times

**Heetec** 3 years ago

A correct - based on comments here

upvoted 1 times

**otsumy** 3 years, 2 months ago

A should be the answer

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 2 times

**ravuri_rk** 3 years, 4 months ago

With managed access schemas, object owners lose the ability to make grant decisions. Only the schema owner (i.e. the role with the OWNERSHIP privilege on the schema) or a role with the MANAGE GRANTS privilege can grant privileges on objects in the schema, including future grants, centralizing privilege management.

upvoted 3 times

**ShagunMittal** 3 years, 7 months ago

Question itself is wrong, privileges are always granted to role, not to users. Users are always granted with roles

upvoted 4 times

**EmiB** 3 years, 7 months ago

In managed access schemas (i.e. schemas created using the CREATE SCHEMA … WITH MANAGED ACCESS syntax), object owners lose the ability to make grant decisions. Only the schema owner (i.e. the role with the OWNERSHIP privilege on the schema) or a role with the global MANAGE GRANTS privilege can grant privileges on objects in the schema.

upvoted 1 times

**lfrad** 3 years, 9 months ago

Only the schema owner (i.e. the role with the OWNERSHIP privilege on the schema) or a role with the MANAGE GRANTS privilege can grant privileges on objects in the schema, including future grants, centralizing privilege management. [for a managed access schema] Here we are however talking about granting permission changes privilege for NEW objects. The schema owner would automatically be granted this privilege on all object within the schema he owns. But a MANAGE GRANTS privileged role could only assign privileges using the "future" keyword

upvoted 2 times

**arpit_dataguy** 3 years, 10 months ago

Answer -- A With managed access schemas, object owners lose the ability to make grant decisions. Only the schema owner (i.e. the role with the OWNERSHIP privilege on the schema) or a role with the MANAGE GRANTS privilege can grant privileges on objects in the schema, including future grants, centralizing privilege management.

upvoted 4 times

**arpit_dataguy** 3 years, 10 months ago

Answer -- A With managed access schemas, object owners lose the ability to make grant decisions. Only the schema owner (i.e. the role with the OWNERSHIP privilege on the schema) or a role with the MANAGE GRANTS privilege can grant privileges on objects in the schema, including future grants, centralizing privilege management.so C is also correct?

upvoted 4 times

**sakis213** 3 years, 10 months ago

so C is also correct?

upvoted 2 times

Question #330 *Topic 1*

What happens when a Data Provider revokes privileges to a share on an object in their source database?

A. The object immediately becomes unavailable for all Data Consumers.

B. Any additional data arriving after this point in time will not be visible to Data Consumers.

C. The Data Consumers stop seeing data updates and become responsible for storage charges for the object.

D. A static copy of the object at the time the privilege was revoked is created in the Data Consumers account.

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**_yyukta** Most Recent 2 years, 2 months ago

A is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**Kvk117** 3 years, 1 month ago

Correct answer

upvoted 2 times

**AravindhTN** 3 years, 4 months ago

Answer correct Revokes access privileges for databases and other supported database objects (schemas, tables, and views) from a share. Revoking privileges on these objects effectively removes the objects from the share, disabling access to the objects granted via the database role in all consumer accounts that have created a database from the share.

upvoted 4 times

Question #331 *Topic 1*

Which command can be used to load data into an internal stage?

A. LOAD

B. COPY

C. GET

D. PUT

**Suggested Answer:** *D*

*Community vote distribution*

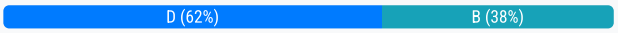


**halol** Highly Voted 3 years, 4 months ago

D is the correct answer

upvoted 14 times

**nexerSnow** Most Recent 1 year, 11 months ago

eiher B or D depending on the source

upvoted 1 times

**_yyukta** 2 years, 2 months ago

D. PUT

upvoted 1 times

**0e504b5** 2 years, 3 months ago

I think it's ambiguous question B or D It depends on the data source COPY unloads the data from table into stage PUT uploads a file from local file system into a stage https://docs.snowflake.com/en/user-guide/data-unload-overview https://docs.snowflake.com/en/sql-reference/sql/put

upvoted 3 times

**Heetec** 2 years, 6 months ago

D correct

upvoted 1 times

**romega2** 2 years, 7 months ago

correct

upvoted 1 times

**Rajivnb** 2 years, 7 months ago

Its not PUT. The question is about loading data (i.e., rows). Not placing FILES. If the questions is about Unloading Data Files then its PUT. Its just Data Load, which means row transfer. So its COPY

upvoted 2 times

**jackjack123** 2 years, 9 months ago

Depending upon from where the data is being loaded.

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**SatyamKishore** 2 years, 11 months ago

Option D as per https://docs.snowflake.com/en/user-guide/data-load-local-file-system-create-stage#types-of-internal-stages

upvoted 1 times

**dmitriypo** 2 years, 11 months ago

B - COPY INTO @stage to unload the data into an internal stage

upvoted 1 times

**learner2023** 3 years ago

PUT command is for uploading the files from local system to cloud environment (snowflake internal stages). GET command is for unload the files from cloud environment (snowflake internal stages) to local system. COPY INTO is used to load data into a Snowflake tableCopy INTO @location is to upload it into a Stage. but this is for external stage

upvoted 1 times

**learner2023** 3 years ago

Selected Answer: D

PUT command is for uploading the files from local system to cloud environment (snowflake internal stages).

GET command is for unload the files from cloud environment (snowflake internal stages) to local system. COPY INTO is used to load data into a Snowflake table

upvoted 1 times

**Rajivnb** 2 years, 9 months ago

Copy INTO @location is to upload it into a Stage. but this is for external stage

upvoted 1 times

**lordchao** 3 years, 1 month ago

chatGPT said it's DChatGPT also says to learn to read documentation.

upvoted 3 times

**lordchao** 3 years, 1 month ago

chatGPT said it's D

upvoted 3 times

**AkexS** 2 years, 9 months ago

ChatGPT also says to learn to read documentation.

upvoted 4 times

**EmiB** 3 years, 1 month ago

D - PUT (load data into a stage (from local)) B - COPY (unload data into a stage (from snowflake))

upvoted 3 times

**ppsingh001** 3 years, 3 months ago

Answer is D We need 'PUT' to uploads data files from a local directory/folder into Internal Stage; Then we can use 'COPY INTO' command to load data into a table from the staged files https://docs.snowflake.com/en/user-guide/data-load-local-file-system-create-stage.html#types-of-internal-stages

upvoted 1 times

**Geera** 3 years, 3 months ago

load/unload -> copy into download/upload -> get/put

upvoted 3 times

**BigDataBB** 3 years, 3 months ago

If the question was "Which command can be used to UNLOAD ..." then is correct copy, but loading on stage means that You are loading external data, so is correct DYou can load data from table to a stage, for instance. its not necessarily external data

upvoted 2 times

**BigDataBB** 3 years, 3 months ago

Selected Answer: D

If the question was "Which command can be used to UNLOAD ..." then is correct copy, but loading on stage means that You are loading external data, so is correct D

upvoted 2 times

**Geera** 3 years, 3 months ago

You can load data from table to a stage, for instance. its not necessarily external data

upvoted 1 times

Question #332 *Topic 1*

What is the MINIMUM Snowflake edition required to use the periodic rekeying of micro-partitions?

A. Enterprise

B. Business Critical

C. Standard

D. Virtual Private Snowflake

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**miadiesel** Highly Voted 3 years, 4 months ago

Enterprise Edition Feature Periodic rekeying requires Enterprise Edition (or higher). To inquire about upgrading, please contact Snowflake Support https://docs.snowflake.com/en/user-guide/security-encryption-manage.html#label-periodic-rekeying

upvoted 8 times

**_yyukta** Most Recent 2 years, 2 months ago

A is correct

upvoted 1 times

**dmitriypo** 2 years, 11 months ago

A is correct

upvoted 1 times

**ViniJsr** 3 years, 2 months ago

A https://docs.snowflake.com/en/user-guide/security-encryption-manage.html#label-periodic-rekeying

upvoted 1 times

**harshagc180** 3 years, 4 months ago

Answer is Enterprise https://docs.snowflake.com/en/user-guide/security-encryption-manage.html#label-periodic-rekeying

upvoted 4 times

Question #333 *Topic 1*

Which stage type can be altered and dropped?

A. Database stage

B. External stage

C. Table stage

D. User stage

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**halol** Highly Voted 3 years, 4 months ago

The answer is correct as we cannot drop the stage associated with a table or user; only named stages (internal or external) can be dropped.

upvoted 6 times

**_yyukta** Most Recent 2 years, 2 months ago

B. External stage

upvoted 1 times

**sandy91** 2 years, 3 months ago

https://docs.snowflake.com/en/sql-reference/sql/drop-stage Internal and external stages both can be drooped.

upvoted 1 times

**metoo23** 2 years, 6 months ago

internal stages are not alterable or droppable, user stages can be altered and dropped, named stages are droppable but not alterable, and table stages are not directly alterable but are dropped when the associated table is dropped

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

B is correct

upvoted 1 times

**learner2023** 3 years ago

https://docs.snowflake.com/en/user-guide/data-load-local-file-system-create-stage

upvoted 1 times

Question #334 *Topic 1*

Which Snowflake object enables loading data from files as soon as they are available in a cloud storage location?

A. Pipe

B. External stage

C. Task

D. Stream

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**halol** Highly Voted 3 years, 4 months ago

I think A

upvoted 8 times

**arpit_dataguy** Highly Voted 3 years, 4 months ago

A. https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html

upvoted 5 times

**_yyukta** Most Recent 2 years, 2 months ago

A. Pipe

upvoted 1 times

**Balathegreat** 2 years, 9 months ago

Snowpipe is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**Harpreetwal** 3 years ago

A is correct

upvoted 1 times

**learner2023** 3 years ago

A is correct answer. Snowpipe enables loading data from files as soon as they're available in a stage. This means you can load data from files in micro-batches, making it available to users within minutes, rather than manually executing COPY statements on a schedule to load larger batches.

upvoted 2 times

**EmiB** 3 years, 1 month ago

A. Pipe

upvoted 1 times

Question #335 *Topic 1*

A user is loading JSON documents composed of a huge array containing multiple records into Snowflake. The user enables the STRIP_OUTER_ARRAY file format option.

What does the STRIP_OUTER_ARRAY file format do?

A. It removes the last element of the outer array.

B. It removes the outer array structure and loads the records into separate table rows.

C. It removes the trailing spaces in the last element of the outer array and loads the records into separate table columns.

D. It removes the NULL elements from the JSON object eliminating invalid data and enables the ability to load the records.

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**AS314** Highly Voted 3 years, 4 months ago

https://cloudyard.in/2021/05/snowflake-strip_outer_array-in-json/

upvoted 7 times

**_yyukta** Most Recent 2 years, 2 months ago

B. It removes the outer array structure and loads the records into separate table rows.

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

STRIP Removes OUTER outer set of square brackets ARRAY []

upvoted 2 times

**SV1122** 3 years, 4 months ago

STRIP_OUTER_ARRAY, Removes the outer set of square brackets [ ] when loading the data, separating the initial array into multiple lines

upvoted 3 times

Question #336 *Topic 1*

Which of the following describes how multiple Snowflake accounts in a single organization relate to various cloud providers?

A. Each Snowflake account can be hosted in a different cloud vendor and region.

B. Each Snowflake account must be hosted in a different cloud vendor and region.

C. All Snowflake accounts must be hosted in the same cloud vendor and region.

D. Each Snowflake account can be hosted in a different cloud vendor, but must be in the same region.

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**halol** Highly Voted 3 years, 4 months ago

I think A

upvoted 9 times

**SnowProCertDec22** Highly Voted 3 years, 4 months ago

A https://docs.snowflake.com/en/user-guide/intro-cloud-platforms.html

upvoted 8 times

**_yyukta** Most Recent 2 years, 2 months ago

A is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

Definitely A

upvoted 1 times

**deadpool1108** 3 years ago

A is correct

upvoted 1 times

**EmiB** 3 years, 1 month ago

A. Each Snowflake account can be hosted in a different cloud vendor and region.

upvoted 2 times

**SV1122** 3 years, 4 months ago

The cloud platform you choose for each Snowflake account is completely independent from your other Snowflake accounts. In fact, you can choose to host each Snowflake account on a different platform, although this may have some impact on data transfer billing when loading data.
https://docs.snowflake.com/en/user-guide/intro-cloud-platforms.html

upvoted 5 times

Question #337 *Topic 1*

If a Snowflake user decides a table should be clustered, what should be used as the cluster key?

A. The columns that are queried in the select clause.

B. The columns with very high cardinality.

C. The columns with many different values.

D. The columns most actively used in the select filters.

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 1 times

**learner2023** 2 years, 6 months ago

D is correct answer. Snowflake recommends prioritizing keys in the order below: Cluster columns that are most actively used in selective filters. For many fact tables involved in date-based queries (for example “WHERE invoice_date > x AND invoice date <= y”), choosing the date column is a good idea. For event tables, event type might be a good choice, if there are a large number of different event types. (If your table has only a small number of different event types, then see the comments on cardinality below before choosing an event column as a clustering key.) If there is room for additional cluster keys, then consider columns frequently used in join predicates, for example “FROM table1 JOIN table2 ON table2.column_A = table1.column_B”.

upvoted 2 times

**halol** 2 years, 10 months ago

D correct

upvoted 4 times

Question #338 *Topic 1*

What are value types that a VARIANT column can store? (Choose two.)

A. STRUCT

B. OBJECT

C. BINARY

D. ARRAY

E. CLOB

**Suggested Answer:** *B, D*

*Community vote distribution*

BD (100%)

**harshagc180** Highly Voted 2 years, 11 months ago

Answer : OBJECT and ARRAY

upvoted 12 times

**SV1122** Highly Voted 2 years, 10 months ago

A VARIANT can store a value of any other type, including OBJECT and ARRAY. The maximum length of a VARIANT is 16 MB.

upvoted 6 times

**Afzy** Most Recent 2 years ago

C and D is correct

upvoted 1 times

**Heetec** 2 years, 1 month ago

B and D

upvoted 1 times

**Amitsnowflake** 2 years, 1 month ago

VARIANT can store a value of any other type, including OBJECT and ARRAY.

upvoted 1 times

**MultiCloudIronMan** 2 years, 3 months ago

Correct

upvoted 1 times

**learner2023** 2 years, 6 months ago

BD is correct. https://docs.snowflake.com/en/sql-reference/data-types-semistructured

upvoted 2 times

**EmiB** 2 years, 8 months ago

OBJECT and ARRAY

upvoted 3 times

**SnowProCertDec22** 2 years, 10 months ago

OBJECT and ARRAY

upvoted 4 times

Question #339 *Topic 1*

A company needs to read multiple terabytes of data for an initial load as part of a Snowflake migration. The company can control the number and size of CSV extract files.

How does Snowflake recommend maximizing the load performance?

A. Use auto-ingest Snowpipes to load large files in a serverless model.

B. Produce the largest files possible, reducing the overall number of files to process.

C. Produce a larger number of smaller files and process the ingestion with size Small virtual warehouses.

D. Use an external tool to issue batched row-by-row inserts within BEGIN TRANSACTION and COMMIT commands.

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**halol** Highly Voted 3 years, 10 months ago

c i think https://www.analytics.today/blog/top-3-snowflake-performance-tuning-tactics#:~:text=Avoid%20Scanning%20Files&text=Before%20copying%20data%2C%20Snowflake%20checks,that%20have%20already%20been%20loaded.

upvoted 7 times

**aemilka** Most Recent 2 years, 1 month ago

Split larger files into a greater number of smaller files to distribute the load among the compute resources in an active warehouse. The number of data files that are processed in parallel is determined by the amount of compute resources in a warehouse. We recommend splitting large files by line to avoid records that span chunks. https://docs.snowflake.com/en/user-guide/data-load-considerations-prepare#:~:text=Split%20larger%20files,that%20span%20chunks.

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

correct

upvoted 1 times

**OTE** 3 years, 8 months ago

I'd go for C. A severless approach (A) is usually not recommended for large files due to the higher costs.

upvoted 3 times

**AS314** 3 years, 10 months ago

https://www.snowflake.com/blog/best-practices-for-data-ingestion/ I think A is correctSnowpipe is designed for continuous ingestion and is built on COPY. The COPY command enables loading batches of data available in external cloud storage or an internal stage. So for the initial stage i think that i s a better solution COPY So the answer is C

upvoted 2 times

**AS314** 3 years, 10 months ago

https://www.snowflake.com/blog/best-practices-for-data-ingestion/ I think A is correct

upvoted 2 times

**BigDataBB** 3 years, 9 months ago

Snowpipe is designed for continuous ingestion and is built on COPY. The COPY command enables loading batches of data available in external cloud storage or an internal stage. So for the initial stage i think that i s a better solution COPY So the answer is C

upvoted 2 times

Question #340 *Topic 1*

For non-materialized views, what column in Information Schema and Account Usage identifies whether a view is secure or not?

A. CHECK_OPTION

B. IS_SECURE

C. IS_UPDATEABLE

D. TABLE_NAME

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 1 times

**examed11** 2 years, 7 months ago

https://docs.snowflake.com/en/user-guide/views-secure#determining-if-a-view-is-secure

upvoted 2 times

**halol** 2 years, 10 months ago

Answer is correct

upvoted 4 times

Question #341 *Topic 1*

The bulk data load history that is available upon completion of the COPY statement is stored where and for how long?

A. In the metadata of the target table for 14 days

B. In the metadata of the pipe for 14 days

C. In the metadata of the target table for 64 days

D. In the metadata of the pipe for 64 days

**Suggested Answer:** *C*

*Community vote distribution*

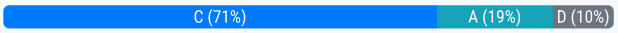


**harshagc180** Highly Voted 2 years, 11 months ago

Correct Answer : C https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html#load-history

upvoted 22 times

**fahfouhi94** Highly Voted 2 years, 9 months ago

data bulk load history ===> in the metadata of the target table for 64 days snowpipe ====> in the metadata of the pipe for 14 days

upvoted 10 times

**PrasantSadangi** Most Recent 1 year, 7 months ago

Load history Bulk data load Stored in the metadata of the target table for 64 days. Available upon completion of the COPY statement as the statement output. Snowpipe Stored in the metadata of the pipe for 14 days. Must be requested from Snowflake via a REST endpoint, SQL table function, or ACCOUNT_USAGE view.

upvoted 1 times

**Amitsnowflake** 2 years, 1 month ago

1. The load history for Snowpipe operations is stored in the metadata of the pipe object , it remains for 14 days. When a pipe is recreatedthe load history is dropped

upvoted 1 times

**singhks** 2 years, 2 months ago

Load history is stored in the metadata of the target table for 64 days

upvoted 1 times

**SA_206** 2 years, 8 months ago

For Snowpipe use copy_history. and for table LOAD_History, both have 14day hisotry https://docs.snowflake.com/en/sql-reference/functions/copy_history.html https://docs.snowflake.com/en/sql-reference/info-schema/load_history.html

upvoted 4 times

**BigDataBB** 2 years, 9 months ago

https://docs.snowflake.com/en/user-guide/data-load-considerations-load.html#load-metadata Snowflake maintains detailed metadata for each table into which data is loaded .... This load metadata expires after 64 days. If the LAST_MODIFIED date for a staged data file is less than or equal to 64 days, the COPY command can determine its load status for a given table and prevent reloading (and data duplication).

upvoted 1 times

**sakis213** 2 years, 10 months ago

It is C

upvoted 2 times

**SV1122** 2 years, 10 months ago

Bulk data load Stored in the metadata of the target table for 64 days. Available upon completion of the COPY statement as the statement output. https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html#load-historySorry the answer is C.

upvoted 2 times

**SV1122** 2 years, 10 months ago

**Selected Answer: D**

Bulk data load Stored in the metadata of the target table for 64 days. Available upon completion of the COPY statement as the statement output.

https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html#load-history

upvoted 2 times

**SV1122** 2 years, 10 months ago

Sorry the answer is C.

upvoted 3 times

**Naveen90** 2 years, 10 months ago

I think, since it is Bulk loading. Answer is A

upvoted 2 times

**halol** 2 years, 10 months ago

A CORRET : https://docs.snowflake.com/en/sql-reference/info-schema/load_history.htmlThis refers to Information schema view. A talks about "metadata of the target table" so this is not correct. The information in target table metadata is retained for 64 days https://docs.snowflake.com/en/user-guide/data-load-considerations-load#load-metadata So, C is correct

upvoted 4 times

**halol** 2 years, 10 months ago

**Selected Answer: A**

A CORRET : https://docs.snowflake.com/en/sql-reference/info-schema/load_history.html

upvoted 4 times

**Def21** 2 years, 4 months ago

This refers to Information schema view. A talks about "metadata of the target table" so this is not correct. The information in target table metadata is retained for 64 days https://docs.snowflake.com/en/user-guide/data-load-considerations-load#load-metadata So, C is correct

upvoted 1 times

Question #342 *Topic 1*

User INQUISITIVE_PERSON has been granted the role DATA_SCIENCE. The role DATA_SCIENCE has privileges OWNERSHIP on the schema MARKETING of the database ANALYTICS_DW.

Which command will show all privileges granted to that schema?

A. SHOW GRANTS ON ROLE DATA_SCIENCE

B. SHOW GRANTS ON SCHEMA ANALYTICS_DW.MARKETING

C. SHOW GRANTS TO USER INQUISITIVE_PERSON

D. SHOW GRANTS OF ROLE DATA_SCIENCE

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Marge23** Most Recent 2 years, 1 month ago

correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 3 months ago

Correct

upvoted 1 times

**examed11** 2 years, 7 months ago

https://docs.snowflake.com/en/sql-reference/sql/show-grants#variants

upvoted 1 times

**sakis213** 2 years, 10 months ago

SHOW GRANTS ON DATABASE, WAREHOUSE, SCHEMA name

upvoted 2 times

Question #343 *Topic 1*

Which of the following are characteristics of security in Snowflake?

A. Account and user authentication is only available with the Snowflake Business Critical edition.

B. Support for HIPAA and GDPR compliance is available for UI Snowflake editions.

C. Periodic rekeying of encrypted data is available with the Snowflake Enterprise edition and higher

D. Private communication to internal stages is allowed in the Snowflake Enterprise edition and higher.

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

Correct

upvoted 1 times

**SV1122** 2 years, 10 months ago

https://docs.snowflake.com/en/user-guide/intro-editions.html

upvoted 3 times

Question #344 *Topic 1*

Which of the following objects can be shared through secure data sharing?

A. Masking policy

B. Stored procedure

C. Task

D. External table

**Suggested Answer:** *D*

*Community vote distribution*

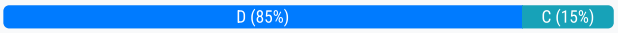


**sakis213** Highly Voted 2 years, 10 months ago

The following Snowflake database objects can be shared: Tables External tables Secure views Secure materialized views Secure UDFs

upvoted 9 times

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

Correct

upvoted 1 times

**SA_206** 2 years, 8 months ago

shared objects are Tables External tables Secure views Secure materialized views Secure UDFs

upvoted 2 times

**SV1122** 2 years, 10 months ago

External Table

upvoted 1 times

Question #345 *Topic 1*

Which formats does Snowflake store unstructured data in? (Choose two.)

A. GeoJSON

B. Array

C. XML

D. Object

E. BLOB

**Suggested Answer:** *B, D*

*Community vote distribution*




**KeshavaMugulur** Highly Voted 2 years, 10 months ago

Answer is B & D

upvoted 15 times

**OTE** Highly Voted 2 years, 8 months ago

Shouldn't the question say SEMI-structured data?

upvoted 10 times

**user_1011** Most Recent 1 year, 5 months ago

The question looks to me incorrect. Shouldnt it say semi structured data

upvoted 1 times

**MultiCloudIronMan** 2 years, 3 months ago

Correct

upvoted 2 times

**psilvabr** 2 years, 7 months ago

Answer is B & D

upvoted 4 times

**leozhang** 2 years, 9 months ago

I am very confusion about this question. Because I read the information that Unstructured data is information that does not fit into a predefined data model or schema. https://docs.snowflake.com/en/user-guide/unstructured-intro.html#sql-functions.

upvoted 2 times

Question #346 *Topic 1*

A user is preparing to load data from an external stage.

Which practice will provide the MOST efficient loading performance?

A. Organize files into logical paths

B. Store the files on the external stage to ensure caching is maintained

C. Use pattern matching for regular expression execution

D. Load the data in one large file

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**SV1122** Highly Voted 2 years, 10 months ago

Both internal (i.e. Snowflake) and external (Amazon S3, Google Cloud Storage, or Microsoft Azure) stage references can include a path (or prefix in AWS terminology). When staging regular data sets, we recommend partitioning the data into logical paths that include identifying details such as geographical location or other source identifiers, along with the date when the data was written. Organizing your data files by path lets you copy any fraction of the partitioned data into Snowflake with a single command. This allows you to execute concurrent COPY statements that match a subset of files, taking advantage of parallel operations. https://docs.snowflake.com/en/user-guide/data-load-considerations-stage.html#organizing-data-by-path

upvoted 5 times

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 1 times

**Rob__C** 2 years, 10 months ago

https://docs.snowflake.com/en/user-guide/data-load-considerations-stage.html#organizing-data-by-path

upvoted 3 times

**paknadeem** 2 years, 10 months ago

The correct answer is A

upvoted 2 times

**halol** 2 years, 10 months ago

The answer is A

upvoted 3 times

Question #347 *Topic 1*

What effect does WAIT_FOR_COMPLETION = TRUE have when running an ALTER WAREHOUSE command and changing the warehouse size?

A. The warehouse size does not change until all queries currently running in the warehouse have completed.

B. The warehouse size does not change until all queries currently in the warehouse queue have completed.

C. The warehouse size does not change until the warehouse is suspended and restarted.

D. It does not return from the command until the warehouse has finished changing its size.

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**SV1122** Highly Voted 2 years, 10 months ago

WAIT_FOR_COMPLETION = FALSE | TRUE When resizing a warehouse, you can use this parameter to block the return of the ALTER WAREHOUSE command until the resize has finished provisioning all its compute resources. Blocking the return of the command when resizing to a larger warehouse serves to notify you that your compute resources have been fully provisioned and the warehouse is now ready to execute queries using all the new resources. https://docs.snowflake.com/en/sql-reference/sql/alter-warehouse.html

upvoted 7 times

**Marge23** Most Recent 2 years, 1 month ago

correct To block the immediate return of the ALTER WAREHOUSE command until the resize is complete, add the WAIT_FOR_COMPLETION parameter. https://docs.snowflake.com/en/sql-reference/sql/alter-warehouse

upvoted 1 times

**MultiCloudIronMan** 2 years, 3 months ago

Correct

upvoted 1 times

Question #348 *Topic 1*

Which of the following can be used when unloading data from Snowflake? (Choose two.)

A. When unloading semi-structured data, it is recommended that the STRIP_OUTER_ARRAY option be used.

B. Use the ENCODING file format option to change the encoding from the default UTF-8.

C. The OBJECT_CONSTRUCT function can be used to convert relational data to semi-structured data.

D. By using the SINGLE = TRUE parameter, a single file up to 5 GB in size can be exported to the storage layer.

E. Use the PARSE_JSON function to ensure structured data will be unloaded into the VARIANT data type.

**Suggested Answer:** *C, D*

*Community vote distribution*




**Marge23** Most Recent 2 years, 1 month ago

correct

upvoted 1 times

**Amitsnowflake** 2 years, 1 month ago

https://docs.snowflake.com/en/user-guide/data-unload-considerations CD

upvoted 1 times

**dmitriypo** 2 years, 5 months ago

C: https://docs.snowflake.com/en/user-guide/data-unload-considerations#unloading-a-relational-table-to-json D: https://docs.snowflake.com/en/user-guide/data-unload-considerations#unloading-to-a-single-file

upvoted 1 times

**SD3** 2 years, 6 months ago

Specifically mentioned in Snowflake Documentation: "Output files are always encoded using UTF-8, regardless of the file format; no other character sets are supported." https://docs.snowflake.com/en/user-guide/intro-summary-unloading#:~:text=File%20encoding,sets%20are%20supported.

upvoted 1 times

**EmiB** 2 years, 7 months ago

D - can be exported to the storage layer - Does storage layer really means STAGE !! ??

upvoted 1 times

**lololi** 2 years, 9 months ago

To me both BCD looks right. Regarding B you can use a file format for unloading (https://docs.snowflake.com/en/user-guide/data-unload-prepare.html) And you can use ENCODING to change it https://docs.snowflake.com/en/sql-reference/sql/create-file-format.html

upvoted 1 times

**sakis213** 2 years, 10 months ago

https://docs.snowflake.com/en/user-guide/data-unload-considerations.html

upvoted 2 times

**SV1122** 2 years, 10 months ago

Answer is CD

upvoted 2 times

Question #349 *Topic 1*

What data is stored in the Snowflake storage layer? (Choose two.)

A. Snowflake parameters

B. Micro-partitions

C. Query history

D. Persisted query results

E. Standard and secure view results

**Suggested Answer:** *B, C*

*Community vote distribution*




**EmiB** Highly Voted 3 years, 1 month ago

B. Micro-partitions C. Query history (also stored in the data layer of SNOWFLAKE.USAGE_SCHEMA)

upvoted 7 times

**PrasantSadangi** Most Recent 1 year, 7 months ago

QUERY_HISTORY view This Account Usage view can be used to query Snowflake query history by various dimensions (time range, session, user, warehouse, etc.) within the last 365 days (1 year).

upvoted 1 times

**levisn1** 1 year, 10 months ago

I think is B and C - in the documentation when it uses the word persisted it talks about 24h period of retention, which is the result set retention time living in the service cloud layer. So i don't think D is correct. https://docs.snowflake.com/en/user-guide/querying-persisted-results

upvoted 1 times

**0e504b5** 2 years, 3 months ago

Service Layer: Which accepts SQL requests from users, coordinates queries, managing transactions and results. Logically, this can be assumed to hold the result cache – a cached copy of the results of every query executed. https://community.snowflake.com/s/article/Caching-in-the-Snowflake-Cloud-Data-Platform D cannot be right because the services layer stores cache, not the the storage layer

upvoted 2 times

**ganeshra** 2 years, 4 months ago

B & D. https://www.snowflake.com/data-cloud-glossary/data-storage/

upvoted 1 times

**SamuelZhang** 2 years, 8 months ago

B. Micro-partitions C. Query history

upvoted 1 times

**singhks** 2 years, 8 months ago

B and C are correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**vini15** 3 years ago

should be BC. cannot be E as non-materialised view results are not stored.

upvoted 1 times

**Kvk117** 3 years, 1 month ago

Verified

upvoted 1 times

**AS314** 3 years, 3 months ago

https://docs.snowflake.com/en/user-guide/querying-persisted-results.html Persisted here means cached, so Persisted Query Results is the query result cache which is in the cloud service layer, not storage layer

upvoted 3 times

Question #350 *Topic 1*

A data provider wants to share data with a consumer who does not have a Snowflake account. The provider creates a reader account for the consumer following these steps:

1. Created a user called "CONSUMER"
2. Created a database to hold the share and an extra-small warehouse to query the data
3. Granted the role PUBLIC the following privileges: Usage on the warehouse, database, and schema, and SELECT on all the objects in the share

Based on this configuration what is true of the reader account?

A. The reader account will automatically use the Standard edition of Snowflake.

B. The reader account compute will be billed to the provider account.

C. The reader account can clone data the provider has shared, but cannot re-share it.

D. The reader account can create a copy of the shared data using CREATE TABLE AS...

**Suggested Answer:** *B*

*Community vote distribution*




**halol** Highly Voted 2 years, 10 months ago

Answer is B: https://docs.snowflake.com/en/user-guide/data-sharing-reader-create.html#overview

upvoted 16 times

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 1 times

**fahfouhi94** 2 years, 9 months ago

the user has not a snowflake account, so the user compute will be billed to provider accoutn

upvoted 4 times

Question #351 *Topic 1*

Which of the following activities consume virtual warehouse credits in the Snowflake environment? (Choose two.)

A. Caching query results

B. Running EXPLAIN and SHOW commands

C. Cloning a database

D. Running a custom query

E. Running COPY commands

**Suggested Answer:** *D, E*

*Community vote distribution*

DE (100%)

**SnowProCertDec22** Highly Voted 2 years, 10 months ago

https://docs.snowflake.com/en/user-guide/warehouses.html

upvoted 8 times

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 1 times

**fahfouhi94** 2 years, 9 months ago

correct answer ===> DE

upvoted 1 times

**SV1122** 2 years, 10 months ago

A warehouse provides the required resources, such as CPU, memory, and temporary storage, to perform the following operations in a Snowflake session: Executing SQL SELECT statements that require compute resources (e.g. retrieving rows from tables and views). Performing DML operations, such as: Updating rows in tables (DELETE , INSERT , UPDATE). Loading data into tables (COPY INTO ). Unloading data from tables (COPY INTO ). Note - To perform these operations, a warehouse must be running and in use for the session. While a warehouse is running, it consumes Snowflake credits.
https://docs.snowflake.com/en/user-guide/warehouses.html

upvoted 4 times

Question #352 *Topic 1*

When loading data into Snowflake, the COPY command supports which of the following?

A. Joins

B. Filters

C. Column reordering

D. Aggregates

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**SnowProCertDec22** Highly Voted 2 years, 4 months ago

https://docs.snowflake.com/en/user-guide/data-load-transform.html#reorder-csv-columns-during-a-load

upvoted 5 times

**MultiCloudIronMan** Most Recent 2 years ago

Verified

upvoted 2 times

Question #353 *Topic 1*

What is cached during a query on a virtual warehouse?

A. All columns in a micro-partition

B. Any columns accessed during the query

C. The columns in the result set of the query

D. All rows accessed during the query

**Suggested Answer:** *B*

*Community vote distribution*




**Lematthew31** 2 years ago

In Snowflake, micro-partitions are the fundamental unit of data storage. They contain a subset of columns from the table and are compressed and encoded for efficient storage and query processing. When a query is executed on a virtual warehouse in Snowflake, only the micro-partitions containing the columns accessed during the query are cached. This approach optimizes caching resources by storing only the necessary data for the query, rather than caching all columns or all rows.

upvoted 4 times

**Sk_3** 2 years, 6 months ago

Its such a confusingquestion- few point. 1its asking about cached in VM.So micro -partion information get stored at cloud layer level.So this out 2. For Data cache it's header and column data are stored on SSD of virtual warehouse. Virtual warehouse first read locally available data (SSD of Virtual warehouse) and then read remaining from remote cloud storage (actual snowflake's storage system). Data keeps dropping on the bases of Least Used fashion when cache storages fills. which means B and D 2 option left.Eliminating logic ...It does not cache all row data ,it caches only those column data which was output of previous query to ... D is out and left with B

upvoted 2 times

**pengrobert** 2 years, 6 months ago

both B and D

upvoted 1 times

**Rajivnb** 2 years, 7 months ago

https://billigence.com/blog/snowflake-data-caching/ . There is no clear explanation from Snowflake on this. But this site tells that the entire Micro Partition accessed is cached. In that case Option A (All Columns from a Micropartition) is correct. "The query will use all of the micro-partitions, which means neither the first or second query are actually pruning any micro-partitions in your table scan. It's important to understand that the cache is made up of the entire micro-partitions that were fetched and not just the records that were selected in the first query"

upvoted 3 times

**jAtlas7** 2 years, 7 months ago

As explained by a few, https://billigence.com/blog/snowflake-data-caching/ has more explanation on than other sites around the levels of caching and what the specific VIRTUAL WAREHOUSE CACHE does - It stores the whole micro partition used in the original query meaning that any other columns in that micro partition would also be stored.

upvoted 1 times

**v1tjkee** 2 years, 7 months ago

Why not C, please?

upvoted 2 times

**KindleAce** 2 years, 8 months ago

Seems D: https://billigence.com/blog/snowflake-data-caching/ Virtual Warehouse Cache This cache is implemented in the compute layer and stores the micro-partitions that were used from the first query and leaves them on your warehouse for future queries. If the next query needs those same micro-partitions to complete its task, it'll use the Warehouse Cache, rather than fetching the micro-partitions again. The query will use all of the micro-partitions, which means neither the first or second query are actually pruning any micro-partitions in your table scan. It's important to understand that the cache is made up of the entire micro-partitions that were fetched and not just the records that were selected in the first query. The Warehouse Cache, in this instance, can optimize the performance of the query

upvoted 2 times

**OTE** 3 years, 1 month ago

B is correct. Each warehouse, when running, maintains a cache of table data accessed as queries are processed by the warehouse. https://docs.snowflake.com/en/user-guide/warehouses-considerations

upvoted 3 times

**Cassielovedata** 3 years, 2 months ago

If I consider the sequence of the data scan: 1. according to the select command, find the specific columns, it means all columns accessed 2.according to the micro-partition metadata, find all micro partition that potentially have the values I needed. it means ALL rows in all the eligible micro-partitions. 3. get the rows I needed as the final results 4. cached it. If it is work in this way, I guess the B is better than D as the final result cache may not include all rows accessed in step 2.

upvoted 4 times

**Cassielovedata** 3 years, 2 months ago

If I consider the sequence of the data scan: 1. according to the select command, find the specific columns, it means all columns accessed 2.according to the micro-partition metadata, find all micro partition that potentially have the values I needed. it means ALL rows in all the eligible micro-partitions. 3. get the rows I needed as the final results 4. cached it. If it is work in this way, I guess the B is better than D as the final result cache may not include all rows accessed in step 2.Then it should be A. All columns in a Micro Partition.

upvoted 4 times

**Rajivnb** 2 years, 7 months ago

Then it should be A. All columns in a Micro Partition.

upvoted 1 times

**lfrad** 3 years, 3 months ago

B “As Snowflake is a columnar data warehouse, it automatically returns the columns needed rather then the entire row to further help maximize query performance” https://community.snowflake.com/s/article/Caching-in-Snowflake-Data-Warehouse

upvoted 3 times

**fahfouhi94** 3 years, 3 months ago

I think D is correct

upvoted 1 times

**subbie** 3 years, 3 months ago

B seems correct

upvoted 1 times

**sailoo** 3 years, 3 months ago

Seems D. https://community.snowflake.com/s/article/Caching-in-Snowflake-Data-Warehouse Local Disk Cache: Which is used to cache data used by SQL queries. Whenever data is needed for a given query it's retrieved from the Remote Disk storage, and cached in SSD and memory.

upvoted 2 times

**tinkofee** 3 years, 4 months ago

Selected Answer: D

Shouldn't it be D?

upvoted 3 times

**tinkofee** 3 years, 4 months ago

Shouldn't it be D?No, because snowflake is column-based storage and can efficiently filter out columns not touched by a query.

upvoted 3 times

**Mitchdu** Most Recent 1 year, 10 months ago

No, because snowflake is column-based storage and can efficiently filter out columns not touched by a query.

upvoted 1 times

Question #354 *Topic 1*

What is the default character set used when loading CSV files into Snowflake?

A. UTF-8

B. UTF-16

C. ISO 8859-1

D. ANSI_X3.4

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**MultiCloudIronMan** Most Recent 2 years ago

Verified

upvoted 2 times

**SV1122** 2 years, 4 months ago

For delimited files (CSV, TSV, etc.), the default character set is UTF-8. https://docs.snowflake.com/en/user-guide/intro-summary-loading.html

upvoted 3 times

Question #355 *Topic 1*

Which of the following describes external functions in Snowflake?

A. They are a type of User-defined Function (UDF).

B. They contain their own SQL code.

C. They call code that is stored inside of Snowflake.

D. They can return multiple rows for each row received.

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**halol** Highly Voted 2 years, 10 months ago

A correct An external function is a type of UDF. Unlike other UDFs, an external function does not contain its own code; instead, the external function calls code that is stored and executed outside Snowflake

upvoted 8 times

**PrasantSadangi** Most Recent 1 year, 7 months ago

https://docs.snowflake.com/en/sql-reference/external-functions-introduction

upvoted 1 times

**MultiCloudIronMan** 2 years, 3 months ago

Correct

upvoted 1 times

**Andywu213** 2 years, 10 months ago

A https://docs.snowflake.com/en/sql-reference/external-functions-introduction.html

upvoted 2 times

Question #356 *Topic 1*

Which of the following are valid methods for authenticating users for access into Snowflake? (Choose three.)

A. SCIM

B. Federated authentication

C. TLS 1.2

D. Key-pair authentication

E. OAuth

F. OCSP authentication

**Suggested Answer:** *B, D, E*

*Community vote distribution*

BDE (100%)

**harshagc180** Highly Voted 2 years, 10 months ago

Sorry, answer is B,D and E https://docs.snowflake.com/en/user-guide/authentication.html

upvoted 16 times

**conrad2023** Most Recent 2 years, 2 months ago

BDE should be correct.

upvoted 1 times

**KnightVictor** 2 years, 5 months ago

for sure

upvoted 1 times

**MultiCloudIronMan** 2 years, 6 months ago

Correct answer

upvoted 1 times

**sailoo** 2 years, 9 months ago

SCIM nothing to do with user authentication so B,D,E

upvoted 2 times

**kokosz** 2 years, 10 months ago

bde is correct

upvoted 2 times

**harshagc180** 2 years, 11 months ago

Correct Answer : A,B & D

upvoted 1 times

Question #357 *Topic 1*

A user has a standard multi-cluster warehouse auto-scaling policy in place.

Which condition will trigger a cluster to shut-down?

A. When after 2-3 consecutive checks the system determines that the load on the most-loaded cluster could be redistributed.

B. When after 5-6 consecutive checks the system determines that the load on the most-loaded cluster could be redistributed.

C. When after 5-6 consecutive checks the system determines that the load on the least-loaded cluster could be redistributed.

D. When after 2-3 consecutive checks the system determines that the load on the least-loaded cluster could be redistributed.

**Suggested Answer:** *D*

*Community vote distribution*

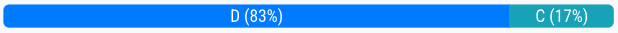


**PrasantSadangi** Most Recent 1 year, 7 months ago

https://hevodata.com/learn/snowflake-scaling-policy/#:~:text=Every%20minute%2C%20the%20system%20checks,1%20every%202%2D3%20minutes.

upvoted 1 times

**MultiCloudIronMan** 1 year, 9 months ago

Selected Answer: C

CD are correct depending on the policy

upvoted 1 times

**MultiCloudIronMan** 1 year, 9 months ago

CD are correct depending on the policyAnswer D bcos it clearly says Standard

upvoted 1 times

**MultiCloudIronMan** 1 year, 9 months ago

Answer D bcos it clearly says Standard

upvoted 2 times

**37ceea2** 1 year, 11 months ago

The right answer is C and D, C is valid for Standard policy, D for Economy.

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

correct

upvoted 1 times

**SV1122** 3 years, 10 months ago

After 2 to 3 consecutive successful checks (performed at 1 minute intervals), which determine whether the load on the least-loaded cluster could be redistributed to the other clusters without spinning up the cluster again. https://docs.snowflake.com/en/user-guide/warehouses-multicluster.html

upvoted 3 times

**AravindhTN** 3 years, 10 months ago

Answer Correct https://docs.snowflake.com/en/user-guide/warehouses-multicluster.html#:~:text=The%20scaling%20policy%20for%20a%20multi%2Dcluster%20warehouse%20only%20applies,or%20shut%20down%20individual%20clusters.&text=Prevents%2Fminimizes%20queuing%20by%20favoring%20starting%20additional%20clusters%20over%20conserving%20credits.

upvoted 4 times

Question #358 *Topic 1*

What is the minimum Snowflake edition needed for database failover and fail-back between Snowflake accounts for business continuity and disaster recovery?

A. Standard

B. Enterprise

C. Business Critical

D. Virtual Private Snowflake

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**harshagc180** Highly Voted 3 years, 4 months ago

answer is C Business Critical https://docs.snowflake.com/en/user-guide/replication-intro.html#business-continuity-and-disaster-recovery

upvoted 12 times

**halol** Highly Voted 4 months ago

C is the answer

upvoted 9 times

**MultiCloudIronMan** 1 year, 9 months ago

Agreed

upvoted 1 times

**fahfouhi94** 4 months ago

Selected Answer: C

C is correct

upvoted 3 times

**_yyukta** 2 years, 2 months ago

C. Business Critical

upvoted 1 times

**sailoo** 4 months ago

https://community.snowflake.com/s/question/0D53r00009kC6WhCAK/can-a-task-have-multiple-sql-queries C is the answer

upvoted 2 times

**baskar14** 2 years, 6 months ago

Business Critical

upvoted 1 times

**_yyukta** 4 months ago

Selected Answer: C

C is correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

Business Critical

upvoted 1 times

**MultiCloudIronMan** Most Recent 4 months ago

Selected Answer: C

correct

upvoted 1 times

**miadiesel** 3 years, 3 months ago

Requires Business Critical (or higher). https://docs.snowflake.com/en/user-guide/database-failover-config.html#failing-over-databases-across-multiple-accounts

upvoted 3 times

Question #359 *Topic 1*

How would a user execute a series of SQL statements using a task?

A. Include the SQL statements in the body of the task CREATE TASK mytask .. AS INSERT INTO target1 SELECT .. FROM stream_s1 WHERE .. INSERT INTO target2 SELECT .. FROM stream_s1 WHERE ..

B. A stored procedure can have only one DML statement per stored procedure invocation and therefore the user should sequence stored procedure calls in the task definition CREATE TASK mytask .... AS call stored_proc1(); call stored_proc2();

C. Use a stored procedure executing multiple SQL statements and invoke the stored procedure from the task. CREATE TASK mytask .... AS call stored_proc_multiple_statements_inside();

D. Create a task for each SQL statement (e.g. resulting in task1, task2, etc.) and string the series of SQL statements by having a control task calling task1, task2, etc. sequentially.

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**halol** Highly Voted 3 years, 4 months ago

C is the answer

upvoted 9 times

**_yyukta** Most Recent 2 years, 2 months ago

C is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

C is correct

upvoted 3 times

**sailoo** 3 years, 3 months ago

https://community.snowflake.com/s/question/0D53r00009kC6WhCAK/can-a-task-have-multiple-sql-queries C is the answer

upvoted 2 times

Question #360 *Topic 1*

How many resource monitors can be assigned at the account level?

A. 1

B. 2

C. 3

D. 4

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**PrasantSadangi** Most Recent 1 year, 7 months ago

An account monitor monitors the credit usage of all the warehouses in the account. An account can only have one account monitor. A warehouse monitor monitors the credit usage of the warehouses assigned to the resource monitor. An account can have multiple warehouse monitors. A warehouse monitor can have one or more warehouses assigned to it, but each warehouse can only be assigned to one resource monitor.

upvoted 1 times

**_yyukta** 2 years, 2 months ago

A is correct

upvoted 1 times

**dmitriypo** 2 years, 11 months ago

A https://docs.snowflake.com/en/user-guide/resource-monitors

upvoted 2 times

**Kelvin_Yu** 2 years, 12 months ago

https://docs.snowflake.com/en/user-guide/resource-monitors#assignment-of-resource-monitors:~:text=A%20single%20monitor%20can%20be%20set%20at%20the%20account%20level%20to%20control%20credit%20usage%20for%20all%20warehouses%20in%20your%20account.

upvoted 1 times

**sachinghalme** 3 years, 2 months ago

A is correct

upvoted 2 times

**fahfouhi94** 3 years, 3 months ago

A is correct

upvoted 2 times

**SV1122** 3 years, 4 months ago

A single monitor can be set at the account level to control credit usage for all warehouses in your account. https://docs.snowflake.com/en/user-guide/resource-monitors.html#monitor-level

upvoted 4 times

**SnowProCertDec22** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/resource-monitors.html

upvoted 2 times

Question #361 *Topic 1*

Data storage for individual tables can be monitored using which commands and/or objects? (Choose two.)

A. SHOW STORAGE BY TABLE;

B. SHOW TABLES;

C. Information Schema -> TABLE_HISTORY

D. Information Schema -> TABLE_FUNCTION

E. Information Schema -> TABLE_STORAGE_METRICS

**Suggested Answer:** *B, E*

*Community vote distribution*

BE (100%)

**SV1122** Highly Voted 3 years, 4 months ago

B. SHOW TABLES E. TABLE_STORAGE_METRICS These two options will show bytes stored. https://docs.snowflake.com/en/sql-reference/sql/show-tables.html https://docs.snowflake.com/en/sql-reference/info-schema/table_storage_metrics.html

upvoted 9 times

**harshagc180** Highly Voted 3 years, 4 months ago

Answer : B and E

upvoted 5 times

**_yyukta** Most Recent 2 years, 2 months ago

B And E are correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**Kvk117** 3 years, 1 month ago

Verified

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

B&E ==> tested

upvoted 1 times

**sailoo** 3 years, 3 months ago

B,E is the right answer

upvoted 1 times

Question #362 *Topic 1*

How would a user run a multi-cluster warehouse in maximized mode?

A. Configure the maximum clusters setting to "Maximum."

B. Turn on the additional clusters manually after starting the warehouse.

C. Set the minimum Clusters and maximum Clusters settings to the same value.

D. Set the minimum clusters and maximum clusters settings to different values.

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**CuriouStudious** Highly Voted 3 years, 5 months ago

It is C. If min=max, there is no room for increasing any clusters and min and max would be same. Hence same value for maximized mode. If min < max then there is room for more clusters to go up till max, hence best for scaling.

upvoted 9 times

**KeshavaMugulur** Highly Voted 3 years, 5 months ago

C is right Min=max = Maximized Max > Min = Auto scale

upvoted 7 times

**_yyukta** Most Recent 2 years, 2 months ago

C is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

for sure, min=max

upvoted 1 times

**ChrisQ** 3 years ago

Should be C

upvoted 2 times

**dementor** 3 years, 2 months ago

C is the right answer

upvoted 2 times

Question #363 *Topic 1*

What internal stages are available in Snowflake? (Choose three.)

A. Schema stage

B. Named stage

C. User stage

D. Stream stage

E. Table stage

F. Database stage

**Suggested Answer:** *B, C, E*

*Community vote distribution*

BCE (100%)

**KeshavaMugulur** Highly Voted 3 years, 5 months ago

It should be BCE User stage, Table Stage and Named stage are internal stages of snowflake

upvoted 12 times

**SnowProCertDec22** Highly Voted 3 years, 4 months ago

https://hevodata.com/learn/snowflake-stages/#ist

upvoted 6 times

**PrashantGupta1616** Most Recent 1 year, 11 months ago

its BCE

upvoted 1 times

**_yyukta** 2 years, 2 months ago

BCE are correct

upvoted 1 times

**conrad2023** 2 years, 8 months ago

See: https://docs.snowflake.com/en/user-guide/data-load-local-file-system-create-stage

upvoted 1 times

**MultiCloudIronMan** 3 years ago

verified

upvoted 1 times

Question #364 *Topic 1*

Which stages are used with the Snowflake PUT command to upload files from a local file system? (Choose three.)

A. Schema Stage

B. User Stage

C. Database Stage

D. Table Stage

E. External Named Stage

F. Internal Named Stage

**Suggested Answer:** *B, D, F*

*Community vote distribution*

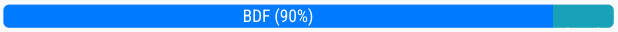


**KeshavaMugulur** Highly Voted 3 years, 5 months ago

BDF is right one. User stage, Table Stage and named internal staged

upvoted 14 times

**_yyukta** Most Recent 2 years, 2 months ago

BDF are correct

upvoted 1 times

**MultiCloudIronMan** 3 years ago

Verified

upvoted 2 times

**GiselaS** 3 years, 1 month ago

BDF Uploads (i.e. stages) data files from a local directory/folder on a client machine to one of the following Snowflake stages: Named internal stage. Internal stage for a specified table. Internal stage for the current user.

upvoted 2 times

**EmiB** 3 years, 1 month ago

PUT does not support uploading files to external stages. To upload files to external stages, use the utilities provided by the cloud service.

upvoted 2 times

**fahfouhi94** 3 years, 3 months ago

https://docs.snowflake.com/en/sql-reference/sql/put.html

upvoted 2 times

**BigDataBB** 3 years, 3 months ago

if i made a PUT i yet haven't a table so i can put on " External Named Stage" or "Internal Named Stage" or "User Stage"

upvoted 1 times

Question #365 *Topic 1*

Which data type can store more than one type of data structure?

A. JSON

B. BINARY

C. VARCHAR

D. VARIANT

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**_yyukta** Most Recent 2 years, 2 months ago

D. VARIANT

upvoted 1 times

**conrad2023** 2 years, 8 months ago

VARIANT

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

variant

upvoted 3 times

Question #366 *Topic 1*

User-level network policies can be created by which of the following roles? (Choose two.)

A. ROLEADMIN

B. ACCOUNTADMIN

C. SYSADMIN

D. SECURITYADNIN

E. USERADMIN

**Suggested Answer:** *B, D*

*Community vote distribution*




**KeshavaMugulur** Highly Voted 3 years, 5 months ago

B & D is right. https://docs.snowflake.com/en/sql-reference/ddl-user-security.html

upvoted 12 times

**_yyukta** Most Recent 2 years, 2 months ago

BD are correct

upvoted 1 times

**MahDeep** 2 years, 7 months ago

Only security administrators (i.e. users with the SECURITYADMIN role) or higher or a role with the global CREATE NETWORK POLICY privilege can create network policies

upvoted 2 times

**dmitriypo** 2 years, 11 months ago

B and D https://docs.snowflake.com/en/user-guide/network-policies#creating-network-policies

upvoted 1 times

**MultiCloudIronMan** 3 years ago

Verified

upvoted 1 times

**SV1122** 3 years, 4 months ago

Only security administrators (i.e. users with the SECURITYADMIN role) or higher or a role with the global CREATE NETWORK POLICY privilege can create network policies. Ownership of a network policy can be transferred to another role. https://docs.snowflake.com/en/user-guide/network-policies.html#:~:text=A%20security%20administrator%20(or%20higher,any%20number%20of%20network%20policies.

upvoted 3 times

Question #367 *Topic 1*

What SQL command would be used to view all roles that were granted to USER1?

A. show grants to user USER1;

B. show grants user USER1;

C. describe user USER1;

D. show grants on user USER1;

**Suggested Answer:** *A*

*Community vote distribution*

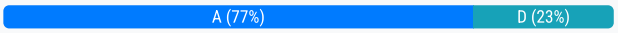


**halol** Highly Voted 3 years, 4 months ago

A https://docs.snowflake.com/en/sql-reference/sql/show-grants.html#variants

upvoted 10 times

**_yyukta** Most Recent 2 years, 2 months ago

A. show grants to user USER1;

upvoted 2 times

**fahfouhi94** 3 years, 3 months ago

A is correct I just tested in my account , command result with TO give info about all role granted to user

upvoted 4 times

**BigDataBB** 3 years, 3 months ago

Selected Answer: D

SHOW GRANTS ON ...

ACCOUNT Lists all the account-level (i.e. global) privileges that have been granted to roles.

object_type object_name Lists all privileges that have been granted on the object.

---------- SHOW GRANTS TO ...

ROLE role_name Lists all privileges and roles granted to the role.

USER user_name Lists all the roles granted to the user. Note that the PUBLIC role, which is automatically available to every user, is not listed.

upvoted 3 times

**BigDataBB** 3 years, 3 months ago

SHOW GRANTS ON ... ACCOUNT Lists all the account-level (i.e. global) privileges that have been granted to roles. object_type object_name Lists all privileges that have been granted on the object. ---------- SHOW GRANTS TO ... ROLE role_name Lists all privileges and roles granted to the role. USER user_name Lists all the roles granted to the user. Note that the PUBLIC role, which is automatically available to every user, is not listed.With your definition it should be A

upvoted 3 times

**MultiCloudIronMan** Highly Voted 2 years, 9 months ago

With your definition it should be A

upvoted 6 times

**SnowProCertDec22** 3 years, 4 months ago

https://docs.snowflake.com/en/sql-reference/sql/show-grants.html#variants SHOW GRANTS TO ROLE role_name Lists all privileges and roles granted to the role.

upvoted 4 times

**KeshavaMugulur** 3 years, 5 months ago

D is right

upvoted 1 times

Question #368 *Topic 1*

Which ACCOUNT_USAGE views are used to evaluate the details of dynamic data masking? (Choose two.)

A. ROLES

B. POLICY_REFERENCES

C. QUERY_HISTORY

D. RESOURCE_MONITORS

E. ACCESS_HISTORY

F. MASKING_POLICIES

**Suggested Answer:** *B, F*

*Community vote distribution*

BF (100%)

**SV1122** Highly Voted 3 years, 4 months ago

Snowflake provides two Account Usage views to obtain information about masking policies: 1. The MASKING POLICIES view provides a list of all masking policies in your Snowflake account. 2. The POLICY_REFERENCES view provides a list of all objects in which a masking policy is set. https://docs.snowflake.com/en/user-guide/security-column-ddm-intro.html

upvoted 11 times

**_yyukta** Most Recent 2 years, 2 months ago

BF are correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**CuriouStudious** 3 years, 4 months ago

B & F. https://docs.snowflake.com/en/user-guide/security-column-ddm-intro.html

upvoted 4 times

**harshagc180** 3 years, 4 months ago

B and F are correct answers

upvoted 4 times

Question #369 *Topic 1*

Which of the following are considerations when using a directory table when working with unstructured data? (Choose two.)

A. A directory table is a separate database object.

B. Directory tables store data file metadata.

C. A directory table will be automatically added to a stage.

D. Directory tables do not have their own grantable privileges.

E. Directory table data can not be refreshed manually.

**Suggested Answer:** *B, D*

*Community vote distribution*

BD (100%)

**_yyukta** Most Recent 2 years, 2 months ago

BD are correct

upvoted 1 times

**dmitriypo** 2 years, 11 months ago

B and D https://docs.snowflake.com/en/user-guide/data-load-dirtables-intro

upvoted 3 times

**KarBiswa** 3 years, 1 month ago

Selected Answer: BD

If three options would have been given then BCD would be the right option but for two options BD are the best selections.

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

If three options would have been given then BCD would be the right option but for two options BD are the best selections.C is wrong. https://docs.snowflake.com/en/user-guide/data-load-dirtables-intro#:~:text=A%20directory%20table%20can%20be%20added%20explicitly%20to%20a%20stage%20when%20the%20stage%20is%20created%20(using%20CREATE%20STAGE)%20or%20later%20(using%20ALTER%20STAGE).

upvoted 1 times

**Kelvin_Yu** 2 years, 12 months ago

C is wrong. https://docs.snowflake.com/en/user-guide/data-load-dirtables-intro#:~:text=A%20directory%20table%20can%20be%20added%20explicitly%20to%20a%20stage%20when%20the%20stage%20is%20created%20(using%20CREATE%20STAGE)%20or%20later%20(using%20ALTER%20STAGE).

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

If three options would have given then BCD are the right options but for two options BD are the best selections.

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

If three options would have given then BCD are the right options but for two options BD are the best selections.https://docs.snowflake.com/en/user-guide/data-load-dirtables#:~:text=it%20is%20an%20implicit%20object%20layered%20on%20a%20stage.%20Both%20external%20(external%20cloud%20storage)%20and%20internal%20(i.e.%20Snowflake)%20stages%20support%20directory%20tables.

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/data-load-dirtables#:~:text=it%20is%20an%20implicit%20object%20layered%20on%20a%20stage.%20Both%20external%20(external%20cloud%20storage)%20and%20internal%20(i.e.%20Snowflake)%20stages%20support%20directory%20tables.

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

correct answer

upvoted 2 times

**Andywu213** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/data-load-dirtables.html

upvoted 1 times

**SV1122** 3 years, 4 months ago

B and D because everything other option is wrong https://docs.snowflake.com/en/user-guide/data-load-dirtables-manage.html

upvoted 3 times

**harshagc180** 3 years, 4 months ago

Answer is correct https://docs.snowflake.com/en/user-guide/data-load-dirtables.html

upvoted 3 times

Question #370 *Topic 1*

The first user assigned to a new account, ACCOUNTADMIN, should create at least one additional user with which administrative privilege?

A. USERADMIN

B. PUBLIC

C. ORGADMIN

D. SYSADMIN

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**halol** Highly Voted 3 years, 10 months ago

A:UserAdmin https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html#using-the-accountadmin-role By default, when your account is provisioned, the first user is assigned the ACCOUNTADMIN role. This user should then create one or more additional users who are assigned the USERADMIN role. All remaining users should be created by the user(s) with the USERADMIN role or another role that is granted the global CREATE USER privilege

upvoted 12 times

**Priyankajade** Most Recent 2 years, 4 months ago

A is correct Answer

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A is correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**dmitriypo** 3 years, 5 months ago

USERADMIN https://docs.snowflake.com/en/user-guide/security-access-control-considerations#using-the-accountadmin-role

upvoted 1 times

**EmiB** 3 years, 8 months ago

A: USERADMIN

upvoted 1 times

**harshagc180** 3 years, 10 months ago

Answer A https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html#

upvoted 4 times

Question #371 *Topic 1*

Which statement describes how Snowflake supports reader accounts?

A. A reader account can consume data from the provider account that created it and combine it with its own data.

B. A consumer needs to become a licensed Snowflake customer as data sharing is only supported between Snowflake accounts.

C. The users in a reader account can query data that has been shared with the reader account and can perform DML tasks.

D. The SHOW MANAGED ACCOUNTS command will view all the reader accounts that have been created for an account.

**Suggested Answer:** *D*

*Community vote distribution*

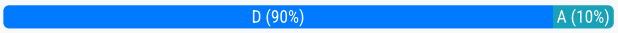


**Vithal_sampagar** Highly Voted 3 years, 10 months ago

D is correct https://docs.snowflake.com/en/sql-reference/sql/show-managed-accounts.html

upvoted 11 times

**KeshavaMugulur** Highly Voted 3 years, 11 months ago

A is right answer

upvoted 5 times

**KeshavaMugulur** 3 years, 11 months ago

A is right answerreader account wont be having their own datawell, it's a bit tricky question! if they have their own data and will able to download the shared data they can combine it?that is right. A is a valid answer

upvoted 5 times

**Madhavan71** 3 years, 10 months ago

reader account wont be having their own data

upvoted 2 times

**Madhavan71** 3 years, 10 months ago

reader account wont be having their own datawell, it's a bit tricky question! if they have their own data and will able to download the shared data they can combine it?that is right. A is a valid answer

upvoted 2 times

**Hopefan** 3 years, 7 months ago

well, it's a bit tricky question! if they have their own data and will able to download the shared data they can combine it?that is right. A is a valid answer

upvoted 2 times

**Hopefan** 3 years, 7 months ago

well, it's a bit tricky question! if they have their own data and will able to download the shared data they can combine it?

upvoted 2 times

**stornati** 3 years, 2 months ago

that is right. A is a valid answer

upvoted 1 times

**37ceea2** Most Recent 1 year, 11 months ago

A and D. What if the reader account will down load that shared data and combine locally with its own data. They can also, apply DML in this case,

upvoted 1 times

**nicholascz** 2 years, 5 months ago

D https://docs.snowflake.com/en/user-guide/data-sharing-reader-create#viewing-reader-accounts

upvoted 1 times

**yaho5** 2 years, 7 months ago

Tricky question. D is for sure a true statement, the command shows all readers accounts, however.. The question is about supporting readers accounts. Does the SHOW MANAGED ACCOUNTS anyhow supports readers accounts? It is just an information. On the other hand, if granted the download privilege, a reader account can download shared data and use other means to compare it with own data. I'm voting for A

upvoted 2 times

**0e504b5** 2 years, 9 months ago

D is correct per the links other uses shared. However, A is not wrong

upvoted 1 times

**Zaaju** 3 years, 2 months ago

D is correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

correct

upvoted 1 times

**dmitriypo** 3 years, 5 months ago

SHOW MANAGED ACCOUNTS https://docs.snowflake.com/en/sql-reference/sql/show-managed-accounts

upvoted 1 times

**chizo** 3 years, 7 months ago

Reader accounts are read only so A is wrong...D is correct

upvoted 1 times

**SV1122** 3 years, 10 months ago

SHOW MANAGED ACCOUNTS Lists the managed accounts created for your account. Currently used by data providers to create reader accounts for their consumers. https://docs.snowflake.com/en/sql-reference/sql/show-managed-accounts.html

upvoted 4 times

Question #372 *Topic 1*

How does Snowflake allow a data provider with an Azure account in central Canada to share data with a data consumer on AWS in Australia?

A. The data provider in Azure Central Canada can create a direct share to AWS Asia Pacific, if they are both in the same organization.

B. The data consumer and data provider can form a Data Exchange within the same organization to create a share from Azure Central Canada to AWS Asia Pacific.

C. The data provider uses the GET DATA workflow in the Snowflake Data Marketplace to create a share between Azure Central Canada and AWS Asia Pacific.

D. The data provider must replicate the database to a secondary account in AWS Asia Pacific within the same organization then create a share to the data consumer's account

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**arpit_dataguy** Highly Voted 2 years, 4 months ago

Answer-- D https://docs.snowflake.com/en/user-guide/secure-data-sharing-across-regions-plaforms.html

upvoted 13 times

**harshagc180** Highly Voted 2 years, 4 months ago

I think answer is D, please confirm

upvoted 10 times

**KnightVictor** Most Recent 1 year, 11 months ago

for sure, replication is the key

upvoted 2 times

**dmitriypo** 1 year, 11 months ago

D is correct

upvoted 1 times

**MultiCloudIronMan** 2 years ago

correct

upvoted 1 times

**fahfouhi94** 2 years, 2 months ago

D is the answer

upvoted 2 times

Question #373 *Topic 1*

Which Snowflake objects can be shared with other Snowflake accounts? (Choose three.)

A. Schemas

B. Roles

C. Secure Views

D. Stored Procedures

E. Tables

F. Secure User-Defined Functions (UDFs)

**Suggested Answer:** *C, E, F*

*Community vote distribution*

CEF (100%)

**harshagc180** Highly Voted 3 years, 10 months ago

Answer : Tables, Secure Views and Secure UDFs https://docs.snowflake.com/en/user-guide/data-sharing-intro.html

upvoted 9 times

**Rameez1** Most Recent 2 years, 4 months ago

Databases Tables Dynamic tables External tables Iceberg tables Secure views Secure materialized views Secure user-defined functions (UDFs)

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

correct

upvoted 1 times

**dmitriypo** 3 years, 5 months ago

CEF https://docs.snowflake.com/en/user-guide/data-sharing-intro

upvoted 1 times

**fahfouhi94** 3 years, 9 months ago

Tables External tables Secure views Secure materialized views Secure UDFs

upvoted 2 times

**SV1122** 3 years, 10 months ago

The following Snowflake database objects can be shared: Tables External tables Secure views Secure materialized views Secure UDFs https://docs.snowflake.com/en/user-guide/data-sharing-intro.html

upvoted 2 times

**kokosz** 3 years, 10 months ago

CEF is the correct answer

upvoted 1 times

**arpit_dataguy** 3 years, 10 months ago

Answer -- CEF

upvoted 1 times

Question #374 *Topic 1*

Which Snowflake feature will allow small volumes of data to continuously load into Snowflake and will incrementally make the data available for analysis?

A. COPY INTO

B. CREATE PIPE

C. INSERT INTO

D. TABLE STREAM

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 5 months ago

PIPE for sure... keyword " continuously"

upvoted 1 times

**fahfouhi94** 2 years, 9 months ago

snwopipe

upvoted 1 times

**harshagc180** 2 years, 10 months ago

Answer is correct CREATE PIPE

upvoted 4 times

Question #375 *Topic 1*

Which Snowflake partner specializes in data catalog solutions?

A. Alation

B. DataRobot

C. dbt

D. Tableau

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**ChrisQ** Highly Voted 3 years ago

To prepare for this kind of questions, we have to remember the entire list of ecosystem? That's so absurd....

upvoted 18 times

**ShawnZh** Most Recent 1 year, 1 month ago

This question looks outupdated, Alation is in Security, Governance & Observability catagory.

upvoted 1 times

**_yyukta** 2 years, 2 months ago

A is correct

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

https://docs.snowflake.com/en/user-guide/ecosystem-all.html

upvoted 3 times

Question #376 *Topic 1*

Which of the following can be executed/called with Snowpipe?

A. A User Defined Function (UDF)

B. A stored procedure

C. A single COPY_INTO statement

D. A single INSERT_INTO statement

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 1 times

**EmiB** 2 years, 8 months ago

C - COPY INTO

upvoted 1 times

**fahfouhi94** 2 years, 9 months ago

I think C https://docs.snowflake.com/en/sql-reference/sql/create-pipe.html

upvoted 2 times

**Jens_B** 2 years, 9 months ago

Shoud be D? https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html#snowpipe-ddl

upvoted 1 times

**Naveen90** 2 years, 10 months ago

C is Correct Answer. Can only COPY_INTO command

upvoted 2 times

**KeshavaMugulur** 2 years, 10 months ago

B is right

upvoted 2 times

Question #377 *Topic 1*

Which snowflake objects will incur both storage and cloud compute charges? (Choose two.)

A. Materialized view

B. Sequence

C. Secure view

D. Transient table

E. Clustered table

**Suggested Answer:** *A, E*

*Community vote distribution*




**SV1122** Highly Voted 2 years, 10 months ago

I think it is A and E. Clustered table needs to undergo clustering as the data changes and Materialized view also undergoes changes every time the underlying data changes or when the view is set to refresh.

upvoted 11 times

**Mitchdu** Most Recent 1 year, 10 months ago

There is no such object type as Clustered table in snowflake, clustering is a feature - not an object type. Thus Materialized view (A) and Transient table (D)

upvoted 2 times

**MultiCloudIronMan** 2 years, 3 months ago

correct

upvoted 1 times

**EmiB** 2 years, 8 months ago

The question doesn't say in which situation incur compute cost. Storage cost - Permanent, Temporary, Transient, MV Compute - creation, maintanance, query ?

upvoted 2 times

**fahfouhi94** 2 years, 9 months ago

normally, transient table also

upvoted 1 times

Question #378 *Topic 1*

What file formats does Snowflake support for loading semi-structured data? (Choose three.)

A. TSV

B. JSON

C. PDF

D. Avro

E. Parquet

F. JPEG

**Suggested Answer:** *B, D, E*

*Community vote distribution*

BDE (100%)

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 1 times

**fahfouhi94** 2 years, 9 months ago

BDE is correct

upvoted 1 times

**SV1122** 2 years, 10 months ago

Snowflake can import semi-structured data from JSON, Avro, ORC, Parquet, and XML formats and store it in Snowflake data types designed specifically to support semi-structured data. And it stores in Array, Object, and Variant formats https://docs.snowflake.com/en/user-guide/semistructured-intro.html#:~:text=Snowflake%20can%20import%20semi%2Dstructured,to%20support%20semi%2Dstructured%20data.

upvoted 2 times

**KeshavaMugulur** 2 years, 10 months ago

DDE is right

upvoted 1 times

**KeshavaMugulur** 2 years, 10 months ago

DDE is rightSorry for typo, it is B D E

upvoted 1 times

**KeshavaMugulur** 2 years, 10 months ago

Sorry for typo, it is B D E

upvoted 4 times

Question #379 *Topic 1*

Which of the following statements about data sharing are true? (Choose two.)

A. New objects created by a Data Provider are automatically shared with existing Data Consumers and Reader Accounts.

B. All database objects can be included in a shared database.

C. Reader Accounts are created by Data Providers.

D. Shared databases are read-only.

E. Reader Accounts are charged for warehouse usage.

**Suggested Answer:** *C, D*

*Community vote distribution*




**harshagc180** Highly Voted 2 years, 10 months ago

Answer : C and D

upvoted 14 times

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

https://www.bing.com/search?q=is+Snowflake+Shared+databases+read-only%3F&qs=n&form=QBRE&sp=-1&lq=0&pq=is+snowflake+shared+databases+read-only%3F&sc=7-40&sk=&cvid=095B1DE142B44CDEBD8765C03FF01A0D&ghsh=0&ghacc=0&ghpl=

upvoted 1 times

**happykeieiekek** 2 years, 3 months ago

CD are correct answer

upvoted 1 times

**OTE** 2 years, 8 months ago

C&D - E is not correct as the data provider is charged for warehouse usage, not the Reader Account (data consumer)

upvoted 2 times

**lfrad** 2 years, 9 months ago

CD right answer

upvoted 1 times

**fahfouhi94** 2 years, 9 months ago

https://docs.snowflake.com/en/user-guide/ui-snowsight-private-sharing-reader-accounts.html

upvoted 2 times

**SalihAlic** 2 years, 9 months ago

Selected Answer: DE

I think it is D and E. Because shared databases are indeed read-only, and reader account are only charged for compute resources/warehouse usage, and not for storage. Data Consumer account don't have to be created by Data Provider necessarily.

upvoted 1 times

**SalihAlic** 2 years, 9 months ago

I think it is D and E. Because shared databases are indeed read-only, and reader account are only charged for compute resources/warehouse usage, and not for storage. Data Consumer account don't have to be created by Data Provider necessarily.Reader Account is not SF's customer, so can't be charged.

upvoted 1 times

**EmiB** 2 years, 8 months ago

Reader Account is not SF's customer, so can't be charged.

upvoted 1 times

Question #380 *Topic 1*

Credit charges for Snowflake virtual warehouses are calculated based on which of the following considerations? (Choose two.)

A. The number of queries executed

B. The number of active users assigned to the warehouse

C. The size of the virtual warehouse

D. The length of time the warehouse is running

E. The duration of the queries that are executed

**Suggested Answer:** *C, D*

*Community vote distribution*

CD (100%)

**Kate0204** Highly Voted 2 years, 11 months ago

CD Snowflake credits are charged based on the number of virtual warehouses you use, how long they run, and their size.

upvoted 8 times

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 2 times

**fahfouhi94** 2 years, 9 months ago

CD is correct answer

upvoted 1 times

**sakis213** 2 years, 10 months ago

Its CD

upvoted 2 times

Question #381 *Topic 1*

Which of the following are handled by the cloud services layer of the Snowflake architecture? (Choose two.)

A. Query execution

B. Data loading

C. Time Travel data

D. Security

E. Authentication and access control

**Suggested Answer:** *D, E*

*Community vote distribution*

DE (100%)

**KeshavaMugulur** Highly Voted 3 years, 5 months ago

D ^ E is right

upvoted 8 times

**SnowProCertDec22** Highly Voted 3 years, 4 months ago

all the rest are not mentioned as a cloud services service https://docs.snowflake.com/en/user-guide/intro-key-concepts.html

upvoted 6 times

**PrashantGupta1616** Most Recent 1 year, 11 months ago

as per chatGPT . Security E. Authentication and access control The cloud services layer is responsible for managing security, authentication, access control, query optimization, metadata management, and other services that coordinate and manage Snowflake's operations.

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 2 times

**fahfouhi94** 3 years, 3 months ago

DE is the answer

upvoted 2 times

**paknadeem** 3 years, 4 months ago

I think A, E

upvoted 1 times

**Kate0204** 3 years, 5 months ago

The cloud services layer is a collection of services that coordinate activities across Snowflake. These services tie together all of the different components of Snowflake in order to process user requests, from login to query dispatch. The cloud services layer also runs on compute instances provisioned by Snowflake from the cloud provider. Services managed in this layer include: Authentication Infrastructure management Metadata management Query parsing and optimization Access control

upvoted 3 times

Question #382 *Topic 1*

What is a responsibility of Snowflake's virtual warehouses?

A. Infrastructure management

B. Metadata management

C. Query execution

D. Query parsing and optimization

E. Permanent storage of micro-partitions

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 2 times

**markpumc** 2 years, 8 months ago

C. Query execution. Snowflake's virtual warehouses are responsible for executing queries and processing data within Snowflake. They are separate from the storage layer, where the actual data is stored in micro-partitions, and the compute layer, where query execution takes place. Virtual warehouses are isolated compute resources that can be scaled up or down as needed to meet changing query performance requirements. When a query is executed, Snowflake automatically selects a virtual warehouse to run the query based on its current workload and available resources. Infrastructure management, metadata management, query parsing and optimization, and permanent storage of micro-partitions are not responsibilities of Snowflake's virtual warehouses. These tasks are performed by other components of the Snowflake architecture.

upvoted 4 times

Question #383 *Topic 1*

What features does Snowflake Time Travel enable?

A. Querying data-related objects that were created within the past 365 days

B. Restoring data-related objects that have been deleted within the past 90 days

C. Conducting point-in-time analysis for BI reporting

D. Analyzing data usage/manipulation over all periods of time

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**KeshavaMugulur** Highly Voted 2 years, 11 months ago

B is right

upvoted 6 times

**Kate0204** Highly Voted 2 years, 11 months ago

Snowflake Time Travel enables accessing historical data (i.e. data that has been changed or deleted) at any point within a defined period. It serves as a powerful tool for performing the following tasks: Restoring data-related objects (tables, schemas, and databases) that might have been accidentally or intentionally deleted. Duplicating and backing up data from key points in the past. Analyzing data usage/manipulation over specified periods of time.

upvoted 5 times

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 1 times

**ssvk2720** 2 years, 6 months ago

No doubt, B is right

upvoted 1 times

**MultiCloudIronMan** 2 years, 6 months ago

There are so many wrong answers

upvoted 1 times

**MultiCloudIronMan** 2 years, 6 months ago

Verified

upvoted 1 times

**EmiB** 2 years, 8 months ago

B - relates only to Ent+ editions C - could be also right in some respects

upvoted 2 times

**fahfouhi94** 2 years, 9 months ago

B is the answer

upvoted 1 times

Question #384 *Topic 1*

Which of the following statements describes a schema in Snowflake?

A. A logical grouping of objects that belongs to a single database

B. A logical grouping of objects that belongs to multiple databases

C. A named Snowflake object that includes all the information required to share a database

D. A uniquely identified Snowflake account within a business entity

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**_yyukta** Most Recent 2 years, 2 months ago

A is correct

upvoted 1 times

**jestem_student** 2 years, 9 months ago

A for sure

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

A for sure

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

A is correct answer

upvoted 1 times

**NikoHacker1995** 3 years, 4 months ago

A is correct

upvoted 1 times

Question #385 *Topic 1*

What is the recommended compressed file size range for continuous data loads using Snowpipe?

A. 8-16 MB

B. 16-24 MB

C. 10-99 MB

D. 100-250 MB

**Suggested Answer:** *D*

*Community vote distribution*

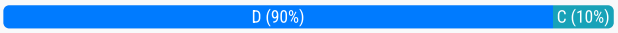


**Kate0204** Highly Voted 3 years, 5 months ago

Snowpipe is typically used to load data that is arriving continuously. File sizing plays an important role in Snowpipe's performance. The recommended file size for data loading is 100-250MB compressed, however, if data is arriving continuously, then try to stage the data within one-minute intervals.

upvoted 6 times

**_yyukta** Most Recent 2 years, 2 months ago

D is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

100-250 MB for sure

upvoted 1 times

**junpal** 3 years, 1 month ago

C is right

upvoted 1 times

**KeshavaMugulur** 3 years, 5 months ago

D is right

upvoted 3 times

Question #386 *Topic 1*

How long is Snowpipe data load history retained?

A. As configured in the CREATE PIPE settings

B. Until the pipe is dropped

C. 64 days

D. 14 days

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**PrasantSadangi** Most Recent 1 year, 7 months ago

When a pipe is paused, event messages received for the pipe enter a limited retention period. The period is 14 days by default. If a pipe is paused for longer than 14 days, it is considered stale.

upvoted 1 times

**_yyukta** 2 years, 2 months ago

B. CREATE PIPE

upvoted 1 times

**dmitriypo** 2 years, 11 months ago

D is correct https://docs.snowflake.com/en/sql-reference/functions/pipe_usage_history

upvoted 1 times

**MultiCloudIronMan** 3 years ago

Correct

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

D is the answer

upvoted 1 times

**BigDataBB** 3 years, 3 months ago

https://docs.snowflake.com/en/sql-reference/functions/pipe_usage_history.html

upvoted 4 times

**NikoHacker1995** 3 years, 4 months ago

Snowpipe Stored in the metadata of the pipe for 14 days. Must be requested from Snowflake via a REST endpoint, SQL table function, or ACCOUNT_USAGE view.

upvoted 3 times

Question #387 *Topic 1*

A company strongly encourages all Snowflake users to self-enroll in Snowflake's default Multi-Factor Authentication (MFA) service to provide increased login security for users connecting to Snowflake.

Which application will the Snowflake users need to install on their devices in order to connect with MFA?

A. Okta Verify

B. Duo Mobile

C. Microsoft Authenticator

D. Google Authenticator

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**SnowProCertDec22** Highly Voted 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/ui-preferences.html#:~:text=Enrolling%20in%20MFA%20(Multi%2DFactor%20Authentication),-MFA%20is%20a&text=This%20second%20form%20of%20authentication,smart%20phone%20or%20similar%20device.

upvoted 5 times

**_yyukta** Most Recent 2 years, 2 months ago

B is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

Mobile

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

Duo for sure

upvoted 1 times

**markpumc** 3 years, 2 months ago

B. Duo Mobile Snowflake integrates with the Duo Mobile app to provide multi-factor authentication (MFA) for users connecting to Snowflake. This means that in order to use MFA, Snowflake users need to have the Duo Mobile app installed on their devices and enroll in Snowflake's MFA service. Okta Verify, Microsoft Authenticator, and Google Authenticator are alternative MFA apps, but they are not directly integrated with Snowflake and may not be supported for use with Snowflake's MFA service.

upvoted 2 times

**halol** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/security-mfa.html

upvoted 4 times

**KeshavaMugulur** 3 years, 5 months ago

it can be C also, Microsoft authenticator to have MFA

upvoted 1 times

Question #388 *Topic 1*

Which URL type allows users to access unstructured data without authenticating into Snowflake or passing an authorization token?

A. Pre-signed URL

B. Scoped URL

C. Signed URL

D. File URL

**Suggested Answer:** *A*

*Community vote distribution*

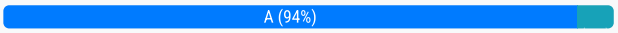


**Kate0204** Highly Voted 3 years, 5 months ago

Pre-signed URLs are used to download or access files, via a web browser for example, without authenticating into Snowflake or passing an authorization token. These URLs are ideal for business intelligence applications or reporting tools that need to display the unstructured file contents.

upvoted 8 times

**niho1993** Most Recent 1 year, 1 month ago

 From Snowflake docs Scoped URL Generated with BUILD_SCOPED_FILE_URL. Works only for unstructured data in internal stages. Purpose: easy, short-lived sharing without Snowflake login or token.  Docs – Scoped URLs Presigned URL Generated with the function GET_PRESIGNED_URL . Works for any files in stages (structured or unstructured). Provides a temporary signed link (like AWS S3 presigned URL). Also does not require Snowflake authentication to access the file.

upvoted 1 times

**_yyukta** 2 years, 2 months ago

A is correct

upvoted 2 times

**MultiCloudIronMan** 3 years ago

Maybe giving wrong answers is part of the strategy here

upvoted 4 times

**KarBiswa** 3 years, 1 month ago

https://docs.snowflake.com/en/user-guide/unstructured-intro

upvoted 1 times

**fahfouhi94** 3 years, 3 months ago

https://docs.snowflake.com/en/sql-reference/functions/get_presigned_url.html

upvoted 2 times

Question #389 *Topic 1*

Where would a Snowflake user find information about query activity from 90 days ago?

A. account_usage.query_history view

B. account_usage.query_history_archive view

C. information_schema.query_history view

D. information_schema.query_history_by_session view

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Kate0204** Highly Voted 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/account-usage/query_history.html

upvoted 7 times

**KeshavaMugulur** Highly Voted 3 years, 11 months ago

Answer: A

upvoted 5 times

**Mallikharjuna452** Most Recent 2 years, 3 months ago

A Is correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

A is correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

Correct

upvoted 1 times

**dmitriypo** 3 years, 5 months ago

https://docs.snowflake.com/en/sql-reference/account-usage/query_history

upvoted 1 times

**fahfouhi94** 3 years, 9 months ago

select * from SNOWFLAKE.ACCOUNT_USAGE.QUERY_HISTORY;

upvoted 1 times

Question #390 *Topic 1*

A marketing co-worker has requested the ability to change a warehouse size on their medium virtual warehouse called MKTG_WH.

Which of the following statements will accommodate this request?

A. ALLOW RESIZE ON WAREHOUSE MKTG_WH TO USER MKTG_LEAD;

B. GRANT MODIFY ON WAREHOUSE MKTG_WH TO ROLE MARKETING;

C. GRANT MODIFY ON WAREHOUSE MKTG_WH TO USER MKTG_LEAD;

D. GRANT OPERATE ON WAREHOUSE MKTG_WH TO ROLE MARKET;

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Kate0204** Highly Voted 3 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/sql/grant-privilege.html

upvoted 9 times

**06c63d0** 2 years, 2 months ago

B and C are quite similar. The difference is that B refers to a role and C refers to a user. Therefore, I think C might be the correct choice, following the principle of least privilege.

upvoted 1 times

**06c63d0** 2 years, 2 months ago

B and C are quite similar. The difference is that B refers to a role and C refers to a user. Therefore, I think C might be the correct choice, following the principle of least privilege.B is correct Privileges are granted to roles, and roles are granted to users, to specify the operations that the users can perform on objects in the system. https://docs.snowflake.com/en/user-guide/security-access-control-privileges

upvoted 1 times

**aemilka** Most Recent 2 years, 1 month ago

B is correct Privileges are granted to roles, and roles are granted to users, to specify the operations that the users can perform on objects in the system. https://docs.snowflake.com/en/user-guide/security-access-control-privileges

upvoted 2 times

**ahadh7621** 2 years, 5 months ago

Answer is B https://docs.snowflake.com/en/user-guide/security-access-control-privileges MODIFY Enables altering any properties of a warehouse, including changing its size. Required to assign a warehouse to a resource monitor. Note that only the ACCOUNTADMIN role can assign warehouses to resource monitors. OPERATE Enables changing the state of a warehouse (stop, start, suspend, resume). In addition, enables viewing current and past queries executed on a warehouse and aborting any executing queries.

upvoted 1 times

**_yyukta** 2 years, 8 months ago

B is correct

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

correct

upvoted 1 times

**fahfouhi94** 3 years, 9 months ago

B is correct answer

upvoted 2 times

**AravindhTN** 3 years, 10 months ago

Selected Answer: B

upvoted 3 times

Question #391 *Topic 1*

Which of the following commands cannot be used within a reader account?

A. CREATE SHARE

B. ALTER WAREHOUSE

C. DROP ROLE

D. SHOW SCHEMAS

E. DESCRIBE TABLE

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Kate0204** Highly Voted 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/data-sharing-reader-create.html#what-is-restricted-allowed-in-a-reader-account

upvoted 9 times

**harshagc180** 4 months ago

Selected Answer: C

Answer : C https://docs.snowflake.com/en/sql-reference/functions/flatten.html

upvoted 5 times

**markpumc** Highly Voted 3 years, 2 months ago

A. CREATE SHARE In Snowflake, a reader account is a special type of user account that has read-only access to data in Snowflake. This means that reader accounts can only perform actions that are related to querying data, such as running SELECT statements and viewing metadata. As a result, reader accounts cannot perform actions that modify the data or metadata stored in Snowflake, such as creating new objects, modifying existing objects, or dropping objects. This includes the CREATE SHARE command, which is used to create a new share and make it available to other users. The other commands listed (ALTER WAREHOUSE, DROP ROLE, SHOW SCHEMAS, and DESCRIBE TABLE) are all allowed within reader accounts and can be used to query metadata and data stored in Snowflake.

upvoted 7 times

**KeshavaMugulur** 4 months ago

It should be C

upvoted 4 times

**_yyukta** 2 years, 2 months ago

A is correct

upvoted 1 times

**kokosz** 4 months ago

Selected Answer: C

C - FLATTEN

upvoted 3 times

**MultiCloudIronMan** 2 years, 9 months ago

Correct

upvoted 1 times

**_yyukta** Most Recent 4 months ago

Selected Answer: C

C is correct

upvoted 1 times

**NikoHacker1995** 3 years, 4 months ago

A is correct (not allowed)

upvoted 2 times

Question #392 *Topic 1*

Which TABLE function helps to convert semi-structured data to a relational representation?

A. CHECK_JSON

B. TO_JSON

C. FLATTEN

D. PARSE_JSON

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**harshagc180** Highly Voted 3 years, 4 months ago

Answer : C https://docs.snowflake.com/en/sql-reference/functions/flatten.html

upvoted 5 times

**_yyukta** Most Recent 2 years, 2 months ago

C is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**jayeshtest** 2 years, 10 months ago

C - FLATTEN

upvoted 1 times

**kokosz** 3 years, 4 months ago

C - FLATTEN

upvoted 3 times

**KeshavaMugulur** 3 years, 5 months ago

It should be C

upvoted 4 times

Question #393 *Topic 1*

Which query profile statistics help determine if efficient pruning is occurring? (Choose two.)

A. Bytes sent over network

B. Percentage scanned from cache

C. Partitions total

D. Bytes spilled to local storage

E. Partitions scanned

**Suggested Answer:** *C, E*

*Community vote distribution*

CE (100%)

**halol** Highly Voted 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/ui-query-profile.html#inefficient-pruning

upvoted 9 times

**Kate0204** Highly Voted 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/ui-query-profile.html#inefficient-pruning

upvoted 6 times

**_yyukta** Most Recent 2 years, 2 months ago

CE are correct

upvoted 1 times

**e_l_i_a_s** 2 years, 3 months ago

The efficiency of pruning can be observed by comparing Partitions scanned and Partitions total statistics in the TableScan operators. If the former is a small fraction of the latter, pruning is efficient. If not, the pruning did not have an effect.

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**NikoHacker1995** 3 years, 4 months ago

C and E are correct, Because, that higher the ratio of existing partitions is to the scanned, that better (more efficient)=

upvoted 4 times

Question #394 *Topic 1*

What are the default Time Travel and Fail-safe retention periods for transient tables?

A. Time Travel - 1 day, Fail-safe - 1 day

B. Time Travel - 0 days, Fail-safe - 1 day

C. Time Travel - 1 day, Failsafe - 0 days

D. Transient tables are retained in neither Fail-safe nor Time Travel.

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**halol** Highly Voted 3 years, 4 months ago

Transient tables can have a Time Travel retention period of either 0 or 1 day. Temporary tables can also have a Time Travel retention period of 0 or 1 day; however, this retention period ends as soon as the table is dropped or the session in which the table was created ends. Transient and temporary tables have no Fail-safe period.

upvoted 5 times

**_yyukta** Most Recent 2 years, 2 months ago

C is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KeshavaMugulur** 3 years, 5 months ago

it should be C

upvoted 4 times

**Kate0204** 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/tables-temp-transient.html#comparison-of-table-types

upvoted 4 times

Question #395 *Topic 1*

Which command is used to unload data from a Snowflake table into a file in a stage?

A. COPY INTO

B. GET

C. WRITE

D. EXTRACT INTO

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**_yyukta** Most Recent 2 years, 2 months ago

A is correct

upvoted 1 times

**MultiCloudIronMan** 3 years ago

Correct

upvoted 1 times

**markpumc** 3 years, 2 months ago

A. COPY INTO

In Snowflake, the "COPY INTO" command is used to unload data from a Snowflake table into a file in a stage. The stage acts as an intermediate storage location for the unloaded data, and the data can then be transferred to an external storage location such as Amazon S3 or Microsoft Azure Blob Storage.

The syntax for the "COPY INTO" command is as follows:

COPY INTO [table_name] FROM [stage_name] FILE_FORMAT = [file_format];

Where [table_name] is the name of the Snowflake table from which data is being unloaded, [stage_name] is the name of the stage where the unloaded data is stored, and [file_format] is the format of the unloaded data files (such as CSV or Parquet).

upvoted 2 times

**markpumc** 3 years, 2 months ago

A. COPY INTO In Snowflake, the "COPY INTO" command is used to unload data from a Snowflake table into a file in a stage. The stage acts as an intermediate storage location for the unloaded data, and the data can then be transferred to an external storage location such as Amazon S3 or Microsoft Azure Blob Storage. The syntax for the "COPY INTO" command is as follows: COPY INTO [table_name] FROM [stage_name] FILE_FORMAT = [file_format]; Where [table_name] is the name of the Snowflake table from which data is being unloaded, [stage_name] is the name of the stage where the unloaded data is stored, and [file_format] is the format of the unloaded data files (such as CSV or Parquet).Example you've given is to extract data from stage to table. Question asked to do vice versa.

upvoted 2 times

**isolated_sniper** 3 years ago

Example you've given is to extract data from stage to table. Question asked to do vice versa.

upvoted 2 times

**SV1122** 3 years, 4 months ago

As illustrated in the diagram below, unloading data to a local file system is performed in two, separate steps: Step 1 Use the COPY INTO command to copy the data from the Snowflake database table into one or more files in a Snowflake stage. In the SQL statement, you specify the stage (named stage or table/user stage) where the files are written. Regardless of the stage you use, this step requires a running, current virtual warehouse for the session if you execute the command manually or within a script. The warehouse provides the compute resources to write rows from the table. Step 2 Use the GET command to download the data files to your local file system. https://docs.snowflake.com/en/user-guide/data-unload-snowflake.html

upvoted 3 times

Question #396 *Topic 1*

What are advantages clones have over tables created with CREATE TABLE AS SELECT statement? (Choose two.)

A. The clone always stays in sync with the original table.

B. The clone has better query performance.

C. The clone is created almost instantly.

D. The clone will have time travel history from the original table.

E. The clone saves space by not duplicating storage.

**Suggested Answer:** *C, E*

*Community vote distribution*




**yaho5** Most Recent 2 years, 1 month ago

CD - correct answers. The question ask about the advantage of clones OVER tables created with a statement CREATE TABLE AS SELECT ... Both tables (clone and the one created with the statement) do not copy the underlaying data, so there is NO advantage of a clone over the other table.

upvoted 1 times

**singhks** 2 years, 8 months ago

C and E are correct.

upvoted 2 times

**MultiCloudIronMan** 3 years ago

Correct

upvoted 1 times

**EmiB** 3 years, 1 month ago

C - Cloning is fast, but not instantaneous, particularly for large objects (e.g. tables). - so ALMOST is correctD - Incorrect - Time travel data is not available for table T1_CLONE. The requested time is either beyond the allowed time travel period or before the object creation time.

upvoted 1 times

**EmiB** 3 years, 1 month ago

Selected Answer: CE

C - Cloning is fast, but not instantaneous, particularly for large objects (e.g. tables). - so ALMOST is correct

upvoted 1 times

**EmiB** 3 years, 1 month ago

D - Incorrect - Time travel data is not available for table T1_CLONE. The requested time is either beyond the allowed time travel period or before the object creation time.

upvoted 1 times

**ShagunMittal** 3 years, 1 month ago

If a table is cloned, historical data for the table clone begins at the time/point when the clone was created. --> this makes option D incorrect for option C, noticing "almost" is important

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

I am going for D&E after going through https://docs.snowflake.com/en/sql-reference/sql/create-clone may be D option is hinting towards creation time of clone it will copy data from time travel

upvoted 1 times

**EmiB** 3 years, 1 month ago

C - Cloning is fast, but not instantaneous, particularly for large objects (e.g. tables).corrected to C&E

upvoted 2 times

**EmiB** 3 years, 1 month ago

Selected Answer: DE

C - Cloning is fast, but not instantaneous, particularly for large objects (e.g. tables).

upvoted 2 times

**EmiB** 3 years, 1 month ago

corrected to C&E

upvoted 1 times

**sakis213** 3 years, 4 months ago

C,E are correct

upvoted 3 times

**JUHYEON** 3 years, 4 months ago

i think C E

upvoted 2 times

**NikoHacker1995** 3 years, 4 months ago

C and E are correct

upvoted 4 times

**halol** 3 years, 4 months ago

https://docs.snowflake.com/en/sql-reference/sql/create-clone.html

upvoted 4 times

Question #397 *Topic 1*

How often are the Account and Table master keys automatically rotated by Snowflake?

A. 30 Days

B. 60 Days

C. 90 Days

D. 365 Days.

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**Kate0204** Highly Voted 3 years, 5 months ago

All Snowflake-managed keys are automatically rotated by Snowflake when they are more than 30 days old.

upvoted 6 times

**_yyukta** Most Recent 2 years, 2 months ago

30 days

upvoted 1 times

**e_l_i_a_s** 2 years, 3 months ago

All Snowflake-managed keys are automatically rotated by Snowflake when they are more than 30 days old. https://docs.snowflake.com/en/user-guide/security-encryption-manage

upvoted 2 times

**singhks** 2 years, 8 months ago

https://docs.snowflake.com/en/user-guide/security-encryption-manage Encryption Key Rotation All Snowflake-managed keys are automatically rotated by Snowflake when they are more than 30 days old. Active keys are retired, and new keys are created. When Snowflake determines the retired key is no longer needed, the key is automatically destroyed. When active, a key is used to encrypt data and is available for usage by the customer. When retired, the key is used solely to decrypt data and is only available for accessing the data.

upvoted 2 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KeshavaMugulur** 3 years, 5 months ago

its 30 Days so A is correct

upvoted 4 times

Question #398 *Topic 1*

Which privilege is required for a role to be able to resume a suspended warehouse if auto-resume is not enabled?

A. USAGE

B. OPERATE

C. MONITOR

D. MODIFY

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**_yyukta** Most Recent 2 years, 2 months ago

B is correct

upvoted 1 times

**e_l_i_a_s** 2 years, 3 months ago

Operate: Grants the ability to start, stop, suspend, or resume a virtual warehouse. Grants the ability to suspend or resume a task. Grants the ability to suspend, resume, or refresh as Dynamic table. Grants the ability to suspend or resume a compute pool. Grants the ability to suspend or resume a Snowpark Container Services service, upgrade service, set, and upset service properties.

upvoted 1 times

**dmitriypo** 2 years, 11 months ago

B is correct https://docs.snowflake.com/en/user-guide/security-access-control-privileges

upvoted 1 times

**MultiCloudIronMan** 3 years ago

correct

upvoted 1 times

**miadiesel** 3 years, 3 months ago

OPERATE: Enables changing the state of a warehouse (stop, start, suspend, resume). In addition, enables viewing current and past queries executed on a warehouse and aborting any executing queries. https://docs.snowflake.com/en/user-guide/security-access-control-privileges.html#virtual-warehouse-privileges

upvoted 1 times

**halol** 3 years, 4 months ago

b correct

upvoted 4 times

Question #399 *Topic 1*

Which statement MOST accurately describes clustering in Snowflake?

A. The database ACCOUNTADMIN must define the clustering methodology for each Snowflake table.

B. Clustering is the way data is grouped together and stored within Snowflake micro-partitions.

C. The clustering key must be included in the COPY command when loading data into Snowflake.

D. Clustering can be disabled within a Snowflake account.

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**_yyukta** Most Recent 2 years, 2 months ago

B is correct

upvoted 1 times

**MultiCloudIronMan** 3 years ago

B is correct

upvoted 1 times

**subbie** 3 years, 4 months ago

B is the most accurate statement about clustering.

upvoted 2 times

**yfswang** 3 years, 4 months ago

Answer should be D Clustering can be disabled within a Snowflake account.Clustering can only be disabled on table level. I would go with B.

upvoted 1 times

**yfswang** 3 years, 4 months ago

Answer should be D Clustering can be disabled within a Snowflake account.

upvoted 1 times

**sakis213** 3 years, 4 months ago

Clustering can only be disabled on table level. I would go with B.

upvoted 1 times

**NikoHacker1995** 3 years, 4 months ago

I guess it is B, but I am happy, if someone can confirm or deny! (P.S. sorry for my bad englisch!)

upvoted 4 times

Question #400 *Topic 1*

Which of the following practices are recommended when creating a user in Snowflake? (Choose two.)

A. Configure the user to be initially disabled.

B. Force an immediate password change.

C. Set a default role for the user.

D. Set the number of minutes to unlock to 15 minutes.

E. Set the user's access to expire within a specified timeframe.

**Suggested Answer:** *B, C*

*Community vote distribution*




**paknadeem** Highly Voted 3 years, 4 months ago

The answer is B,C https://docs.snowflake.com/en/user-guide/admin-user-management.html#creating-users

upvoted 9 times

**_yyukta** Most Recent 2 years, 2 months ago

BC are correct

upvoted 1 times

**e_l_i_a_s** 2 years, 3 months ago

Leave the Force user to change password on first time login checkbox selected to force the user to change their password when they sign in.

upvoted 1 times

**ssvk2720** 3 years ago

B & C are correct answer

upvoted 2 times

**AlexbDku** 3 years, 4 months ago

B&C https://docs.snowflake.com/en/user-guide/admin-user-management.html#creating-users

upvoted 3 times

**Kate0204** 3 years, 5 months ago

B and not sure what else https://docs.snowflake.com/en/user-guide/admin-user-management.html#best-practices-for-password-policies-and-passwords

upvoted 2 times

Question #401 *Topic 1*

Network policies can be applied to which of the following Snowflake objects? (Choose two.)

A. Roles

B. Databases

C. Warehouses

D. Users

E. Accounts

**Suggested Answer:** *D, E*

*Community vote distribution*

DE (100%)

**_yyukta** Most Recent 2 years, 2 months ago

DE are correct

upvoted 2 times

**e_l_i_a_s** 2 years, 3 months ago

Network policies applied to: Account Security Integration User

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**AlexbDku** 3 years, 4 months ago

Users & accounts https://docs.snowflake.com/en/user-guide/network-policies.html#label-verify-network-policy-activation

upvoted 1 times

**miadiesel** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/network-policies.html#label-verify-network-policy-activation

upvoted 3 times

**Kate0204** 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/network-policies.html#label-verify-network-policy-activation

upvoted 4 times

Question #402 *Topic 1*

Where is Snowflake metadata stored?

A. Within the data files

B. In the virtual warehouse layer

C. In the cloud services layer

D. In the remote storage layer

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**_yyukta** Most Recent 2 years, 2 months ago

C is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

Cloud Services Layer

upvoted 1 times

**AlexbDku** 3 years, 4 months ago

Cloud services layer https://docs.snowflake.com/en/user-guide/intro-key-concepts.html#cloud-services

upvoted 2 times

**miadiesel** 3 years, 4 months ago

Metadata management https://docs.snowflake.com/en/user-guide/intro-key-concepts.html#cloud-services

upvoted 4 times

**Kate0204** 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/intro-key-concepts.html#cloud-services

upvoted 4 times

Question #403 *Topic 1*

What columns are returned when performing a FLATTEN command on semi-structured data? (Choose two.)

A. KEY

B. NODE

C. VALUE

D. LEVEL

E. ROOT

**Suggested Answer:** *A, C*

*Community vote distribution*

AC (100%)

**Kate0204** Highly Voted 3 years, 5 months ago

https://docs.snowflake.com/en/sql-reference/functions/flatten.html#output

upvoted 7 times

**PrasantSadangi** Most Recent 1 year, 7 months ago

Output The returned rows consist of a fixed set of columns: +-----+------+------+-------+-------+------+ | SEQ | KEY | PATH | INDEX | VALUE | THIS | |-----+------+------+-------+-------+------| SEQ A unique sequence number associated with the input record; the sequence is not guaranteed to be gap-free or ordered in any particular way. KEY For maps or objects, this column contains the key to the exploded value. PATH The path to the element within a data structure that needs to be flattened. INDEX The index of the element, if it is an array; otherwise NULL. VALUE The value of the element of the flattened array/object. THIS The element being flattened (useful in recursive flattening).

upvoted 1 times

**_yyukta** 2 years, 2 months ago

Ac are correct

upvoted 1 times

**MultiCloudIronMan** 3 years ago

Correct

upvoted 1 times

**miadiesel** 3 years, 4 months ago

https://docs.snowflake.com/en/sql-reference/functions/flatten.html#output

upvoted 3 times

Question #404 *Topic 1*

Which of the following Snowflake features provide continuous data protection automatically? (Choose two.)

A. Internal stages

B. Incremental backups

C. Time Travel

D. Zero-copy clones

E. Fail-safe

**Suggested Answer:** *C, E*

*Community vote distribution*

CE (100%)

**Kate0204** Highly Voted 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/data-cdp.html

upvoted 6 times

**_yyukta** Most Recent 2 years, 2 months ago

CE are correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**AlexbDku** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/data-cdp.html

upvoted 2 times

**miadiesel** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/data-cdp.html

upvoted 3 times

Question #405 *Topic 1*

A developer is granted ownership of a table that has a masking policy. The developer's role is not able to see the masked data.

Will the developer be able to modify the table to read the masked data?

A. Yes, because a table owner has full control and can unset masking policies.

B. Yes, because masking policies only apply to cloned tables.

C. No, because masking policies must always reference specific access roles.

D. No, because ownership of a table does not include the ability to change masking policies.

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**miadiesel** Highly Voted 3 years, 10 months ago

Object owners (i.e. the role that has the OWNERSHIP privilege on the object) do not have the privilege to unset masking policies. Object owners cannot view column data in which a masking policy applies. https://docs.snowflake.com/en/user-guide/security-column-intro.html#what-are-masking-policies

upvoted 5 times

**Kate0204** Highly Voted 4 months ago

Selected Answer: C

C. Use the standard warehouse scaling policy

upvoted 6 times

**PrasantSadangi** 1 year, 7 months ago

https://docs.snowflake.com/en/user-guide/security-column-intro#what-are-masking-policies Object owners (i.e. the role that has the OWNERSHIP privilege on the object) do not have the privilege to unset masking policies. Object owners cannot view column data in which a masking policy applies.

upvoted 1 times

**SV1122** 4 months ago

Selected Answer: C

The economy policy doesn't start immediately Only if the system estimates there's enough query load to keep the cluster busy for at least 6 minutes. https://docs.snowflake.com/en/user-guide/warehouses-multicluster.html

upvoted 2 times

**d22770a** 2 years ago

The question here is Will the developer be able to modify the table to read the masked data? It is about modifying the table to read the data from masked column. Owner of the table can definitely UNSET the masking policy by using this command ALTER TABLE modify column unset masking policy;

upvoted 1 times

**_yyukta** 4 months ago

Selected Answer: C

C is correct

upvoted 1 times

**_yyukta** 2 years, 8 months ago

D is correct

upvoted 1 times

**MultiCloudIronMan** Most Recent 4 months ago

Selected Answer: C

correct

upvoted 1 times

**AlexbDku** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/security-column-intro.html#what-are-masking-policies:~:text=Object%20owners%20(i.e.%20the%20role%20that%20has%20the%20OWNERSHIP%20privilege%20on%20the%20object)%20do%20not%20have%20the%20privilege%20to%20unset%20masking%20policies

upvoted 2 times

**halol** 3 years, 10 months ago

Object owners (i.e. the role that has the OWNERSHIP privilege on the object) do not have the privilege to unset masking policies. Object owners cannot view column data in which a masking policy applies. https://docs.snowflake.com/en/user-guide/security-column-intro.html#what-are-masking-policies

upvoted 3 times

Question #406 *Topic 1*

How should a virtual warehouse be configured if a user wants to ensure that additional multi-clusters are resumed with no delay?

A. Configure the warehouse to a size larger than generally required

B. Set the minimum and maximum clusters to autoscale

C. Use the standard warehouse scaling policy

D. Use the economy warehouse scaling policy

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**Kate0204** Highly Voted 3 years, 5 months ago

C. Use the standard warehouse scaling policy

upvoted 6 times

**_yyukta** Most Recent 2 years, 2 months ago

C is correct

upvoted 1 times

**MultiCloudIronMan** 3 years ago

correct

upvoted 1 times

**SV1122** 3 years, 4 months ago

The economy policy doesn't start immediately Only if the system estimates there's enough query load to keep the cluster busy for at least 6 minutes.
https://docs.snowflake.com/en/user-guide/warehouses-multicluster.html

upvoted 2 times

Question #407 *Topic 1*

During periods of warehouse contention, which parameter controls the maximum length of time a warehouse will hold a query for processing?

A. STATEMENT_TIMEOUT_IN_SECONDS

B. STATEMENT_QUEUED_TIMEOUT_IN_SECONDS

C. MAX_CONCURRENCY_LEVEL

D. QUERY_TIMEOUT_IN_SECONDS

**Suggested Answer:** *B*

*Community vote distribution*

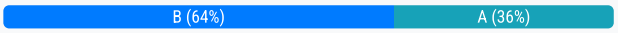


**isolated_sniper** Highly Voted 3 years ago

Answer B STATEMENT_QUEUED_TIMEOUT_IN_SECONDS is a configuration parameter in Snowflake that specifies the maximum amount of time, in seconds, that a statement can remain in a queued state before timing out. STATEMENT_TIMEOUT_IN_SECONDS is a parameter in Snowflake that specifies the maximum amount of time that a query can execute before being automatically terminated by the system. The default value for this parameter is 0, which means there is no timeout limit set.STATEMENT_TIMEOUT_IN_SECONDS - default is 2 days.
https://docs.snowflake.com/en/sql-reference/parameters#statement-timeout-in-seconds

upvoted 6 times

**harshagc180** Highly Voted 4 months ago

Answer - ACCOUNTADMIN https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#

upvoted 8 times

**0e504b5** 2 years, 2 months ago

STATEMENT_TIMEOUT_IN_SECONDS - default is 2 days. https://docs.snowflake.com/en/sql-reference/parameters#statement-timeout-in-seconds

upvoted 2 times

**Nitin001** 4 months ago

Selected Answer: C

ORGADMIN is NOT part of Snowflake Role Hierarchy. So, answer should be C: ACCOUNTADMIN.

upvoted 2 times

**Kate0204** 3 years, 5 months ago

https://docs.snowflake.com/en/sql-reference/parameters.html#statement-queued-timeout-in-seconds

upvoted 6 times

**J** 4 months ago

Selected Answer: B

Org admin super power

upvoted 1 times

**ShawnZh** 1 year, 1 month ago

STATEMENT_TIMEOUT_IN_SECONDS Type Session and Object (for warehouses) Can be set for Account » User » Session; can also be set for individual warehouses Data Type Number Description Amount of time, in seconds, after which a running SQL statement (query, DDL, DML, and so on) is canceled by the system.Key word is RUNNING, the question is for a query of processing. B is for query in the QUEUE.

upvoted 1 times

**lobr** 4 months ago

Selected Answer: C

C is correct

upvoted 1 times

**ShawnZh** 1 year, 1 month ago

Key word is RUNNING, the question is for a query of processing. B is for query in the QUEUE.

upvoted 1 times

**aws_god** Most Recent 4 months ago

Selected Answer: C

https://docs.snowflake.com/en/user-guide/security-access-control-overview#role-hierarchy-and-privilege-inheritance

upvoted 1 times

**PrasantSadangi** 1 year, 7 months ago

STATEMENT_QUEUED_TIMEOUT_IN_SECONDS Type Session and Object (for warehouses) Can be set for Account » User » Session; can also be set for individual warehouses Data Type Number Description Amount of time, in seconds, a SQL statement (query, DDL, DML, etc.) remains queued for a warehouse before it is canceled by the system.

upvoted 1 times

**_yyukta** 2 years, 2 months ago

B is correct

upvoted 1 times

**dmitriypo** 2 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/parameters#statement-timeout-in-seconds

upvoted 1 times

**drserguey** 3 years ago

STATEMENT_QUEUED_TIMEOUT_IN_SECONDS Amount of time, in seconds, a SQL statement (query, DDL, DML, etc.) remains queued for a warehouse before it is canceled by the system. https://docs.snowflake.com/en/sql-reference/parameters#statement-queued-timeout-in-seconds

upvoted 1 times

**MultiCloudIronMan** 3 years ago

VerifiedChange my mind to B

upvoted 1 times

**MultiCloudIronMan** 2 years, 10 months ago

Change my mind to B

upvoted 1 times

**Kvk117** 3 years, 1 month ago

Verified

upvoted 2 times

**SV1122** 3 years, 4 months ago

Amount of time, in seconds, a SQL statement (query, DDL, DML, etc.) remains queued for a warehouse before it is canceled by the system. This parameter can be used in conjunction with the MAX_CONCURRENCY_LEVEL parameter to ensure a warehouse is never backlogged https://docs.snowflake.com/en/sql-reference/parameters.html#statement-queued-timeout-in-seconds

upvoted 3 times

Question #408 *Topic 1*

Files have been uploaded to a Snowflake internal stage. The files now need to be deleted.

Which SQL command should be used to delete the files?

A. PURGE

B. MODIFY

C. REMOVE

D. DELETE

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**Kate0204** Highly Voted 3 years, 5 months ago

https://docs.snowflake.com/en/sql-reference/sql/remove.html

upvoted 7 times

**harshagc180** Highly Voted 4 months ago

Answer - ACCOUNTADMIN https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#

upvoted 8 times

**PrasantSadangi** 1 year, 7 months ago

Staged files can be deleted from a Snowflake stage (user stage, table stage, or named stage) using the following methods: Files that were loaded successfully can be deleted from the stage during a load by specifying the PURGE copy option in the COPY INTO command. After the load completes, use the REMOVE command to remove the files in the stage.

upvoted 1 times

**Nitin001** 4 months ago

Selected Answer: C

ORGADMIN is NOT part of Snowflake Role Hierarchy. So, answer should be C: ACCOUNTADMIN.

upvoted 2 times

**_yyukta** 2 years, 2 months ago

C is correct

upvoted 1 times

**J** 4 months ago

Selected Answer: B

Org admin super power

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**lobr** Most Recent 4 months ago

Selected Answer: C

C is correct

upvoted 1 times

**miadiesel** 3 years, 4 months ago

https://docs.snowflake.com/en/sql-reference/sql/remove.html

upvoted 3 times

**aws_god** 4 months ago

Selected Answer: C

https://docs.snowflake.com/en/user-guide/security-access-control-overview#role-hierarchy-and-privilege-inheritance

upvoted 1 times

**KeshavaMugulur** 3 years, 5 months ago

C is right

upvoted 3 times

Question #409 *Topic 1*

In a Snowflake role hierarchy, what is the top-level role?

A. SYSADMIN

B. ORGADMIN

C. ACCOUNTADMIN

D. SECURITYADMIN

**Suggested Answer:** *C*

*Community vote distribution*

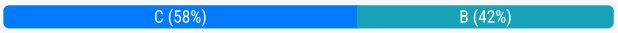


**harshagc180** Highly Voted 3 years, 4 months ago

Answer - ACCOUNTADMIN https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#

upvoted 8 times

**J** Most Recent 1 year, 2 months ago

Org admin super power

upvoted 1 times

**lobr** 1 year, 5 months ago

C is correct

upvoted 1 times

**aws_god** 2 years ago

https://docs.snowflake.com/en/user-guide/security-access-control-overview#role-hierarchy-and-privilege-inheritance

upvoted 1 times

**xardas86** 2 years, 1 month ago

correct

upvoted 1 times

**_yyukta** 2 years, 2 months ago

correct

upvoted 1 times

**SCube** 2 years, 6 months ago

C ORGADMIN is a separate system role that manages operations at the organization level. This role is not included in the hierarchy of system roles.

upvoted 1 times

**kayuhanbiduk** 2 years, 6 months ago

C. ORGADMIN is not part of Role Hierarchy and Privilege Inheritance

upvoted 1 times

**nonamae** 2 years, 6 months ago

AccountAdmin - Role that encapsulates the SYSADMIN and SECURITYADMIN system-defined roles. It is the top-level role in the system and should be granted only to a limited/controlled number of users in your account.

upvoted 1 times

**Amitsnowflake** 2 years, 7 months ago

ACCOUNTADMIN is the top-level role in the system.

upvoted 1 times

**singhks** 2 years, 8 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-overview#label-role-hierarchy-and-privilege-inheritance ORGADMIN is a separate system role that manages operations at the organization level. This role is not included in the hierarchy of system roles.

upvoted 1 times

**MultiCloudIronMan** 2 years, 10 months ago

https://www.bing.com/search?q=In+a+Snowflake+role+hierarchy%2C+what+is+the+top-level+role%3F&aqs=edge..69i57j69i11004&FORM=ANCMS9&PC=U531

upvoted 1 times

**Nitin001** 2 years, 11 months ago

ORGADMIN is NOT part of Snowflake Role Hierarchy. So, answer should be C: ACCOUNTADMIN.

upvoted 2 times

**KnightVictor** 2 years, 11 months ago

Orgadmin

upvoted 1 times

**Hopefan** 2 years, 11 months ago

account admin https://www.google.com/search?q=In+a+Snowflake+role+hierarchy,+what+is+the+top-level+role%3F&rlz=1C1GCEB_enAU913AU913&sxsrf=APwXEdcjcfcT1SmOreRdY6m3KgT2ZqYGtA:1684463664315&source=lnms&tbm=isch&sa=X&ved=2ahUKEwjCx-SSrID_AhUPmVYBHUfeBAYQ_AUoAXoECAEQAw&biw=1920&bih=969&dpr=1#imgrc=7vBSyKRemBqBfM

upvoted 1 times

**bibisa** 3 years ago

Note ORGADMIN is a separate system role that manages operations at the organization level. This role is not included in the hierarchy of system roles.

upvoted 1 times

**MultiCloudIronMan** 3 years ago

**Selected Answer: B**

B is the answer

upvoted 1 times

**MultiCloudIronMan** 3 years ago

B is the answerChanged my answer to C https://www.bing.com/search?q=In+a+Snowflake+role+hierarchy%2C+what+is+the+top-level+role%3F&aqs=edge..69i57j69i11004&FORM=ANCMS9&PC=U531

upvoted 1 times

**MultiCloudIronMan** 2 years, 10 months ago

Changed my answer to C https://www.bing.com/search?q=In+a+Snowflake+role+hierarchy%2C+what+is+the+top-level+role%3F&aqs=edge..69i57j69i11004&FORM=ANCMS9&PC=U531

upvoted 1 times

Question #410 *Topic 1*

By default, which Snowflake role is required to create a share?

A. ORGADMIN

B. SECURITYADMIN

C. SHAREADMIN

D. ACCOUNTADMIN

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**Kate0204** Highly Voted 3 years, 5 months ago

https://docs.snowflake.com/en/sql-reference/sql/create-share.html#access-control-requirements

upvoted 8 times

**_yyukta** Most Recent 2 years, 2 months ago

D is correct

upvoted 1 times

**Avi0407** 2 years, 6 months ago

CREATE SHARE Account Only the ACCOUNTADMIN role has this privilege by default. The privilege can be granted to additional roles as needed.

upvoted 1 times

**MultiCloudIronMan** 3 years ago

Answer is D

upvoted 1 times

**miadiesel** 3 years, 4 months ago

CREATE SHARE: Account :Only the ACCOUNTADMIN role has this privilege by default. The privilege can be granted to additional roles as needed.

https://docs.snowflake.com/en/sql-reference/sql/create-share.html#access-control-requirements

upvoted 2 times

Question #411 *Topic 1*

What happens to historical data when the retention period for an object ends?

A. The data is cloned into a historical object.

B. The data moves to Fail-safe

C. Time Travel on the historical data is dropped.

D. The object containing the historical data is dropped.

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**_yyukta** Most Recent 2 years, 2 months ago

B is correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**miadiesel** 3 years, 4 months ago

When the retention period ends for an object, the historical data is moved into Snowflake Fail-safe: https://docs.snowflake.com/en/user-guide/data-time-travel.html#data-retention-period

upvoted 3 times

**halol** 3 years, 4 months ago

When the retention period ends for an object, the historical data is moved into Snowflake Fail-safe: Historical data is no longer available for querying.

upvoted 3 times

Question #412 *Topic 1*

A company's security audit requires generating a report listing all Snowflake logins (e.g., date and user) within the last 90 days.

Which of the following statements will return the required information?

A. SELECT LAST_SUCCESS_LOGIN, LOGIN_NAME FROM ACCOUNT_USAGE.USERS;

B. SELECT EVENT_TIMESTAMP, USER_NAME FROM table(information_schema.login_history_by_user())

C. SELECT EVENT_TIMESTAMP, USER_NAME FROM ACCOUNT_USAGE.ACCESS_HISTORY;

D. SELECT EVENT_TIMESTAMP, USER_NAME FROM ACCOUNT_USAGE.LOGIN_HISTORY;

**Suggested Answer:** *D*

*Community vote distribution*

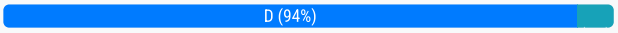


**Kate0204** Highly Voted 3 years, 5 months ago

correct

upvoted 8 times

**_yyukta** Most Recent 2 years, 2 months ago

D. SELECT EVENT_TIMESTAMP, USER_NAME FROM ACCOUNT_USAGE.LOGIN_HISTORY is correct

upvoted 1 times

**MultiCloudIronMan** 3 years ago

Verified

upvoted 1 times

**ShagunMittal** 3 years, 1 month ago

B will be correct only if information needs to be fetched within 7 days

upvoted 2 times

**OTE** 3 years, 1 month ago

D - Correct https://docs.snowflake.com/en/sql-reference/account-usage/login_history B - Incorrect as login_history_by_user function returns login activity within the last 7 days only. https://docs.snowflake.com/en/sql-reference/functions/login_history

upvoted 3 times

**markpumc** 3 years, 2 months ago

B. SELECT EVENT_TIMESTAMP, USER_NAME FROM table(information_schema.login_history_by_user())

The statement in option B will return the required information of the Snowflake logins (date and user) within the last 90 days.

The information_schema.login_history_by_user function returns information about logins to the account by individual users. The returned columns include EVENT_TIMESTAMP and USER_NAME, which can be used to generate the report listing all Snowflake logins within the last 90 days by filtering the results to only show logins within the specified timeframe.

upvoted 2 times

**markpumc** 3 years, 2 months ago

B. SELECT EVENT_TIMESTAMP, USER_NAME FROM table(information_schema.login_history_by_user()) The statement in option B will return the required information of the Snowflake logins (date and user) within the last 90 days. The information_schema.login_history_by_user function returns information about logins to the account by individual users. The returned columns include EVENT_TIMESTAMP and USER_NAME, which can be used to generate the report listing all Snowflake logins within the last 90 days by filtering the results to only show logins within the specified timeframe.The function return login activity only within the last 7 days. So, for 8 to 365 days you'll need to reach the account usage login history view https://docs.snowflake.com/en/sql-reference/functions/login_history

upvoted 2 times

**ganeshra** 2 years, 4 months ago

The function return login activity only within the last 7 days. So, for 8 to 365 days you'll need to reach the account usage login history view https://docs.snowflake.com/en/sql-reference/functions/login_history

upvoted 1 times

**AndroJS** 3 years, 3 months ago

It needs to be in a predicate https://docs.snowflake.com/en/sql-reference/functions/login_history.htmlOnly 7 days so B is not right.

upvoted 1 times

**AndroJS** 3 years, 3 months ago

**Selected Answer: B**

It needs to be in a predicate https://docs.snowflake.com/en/sql-reference/functions/login_history.html

upvoted 1 times

**nzpentest** 3 years, 3 months ago

Only 7 days so B is not right.

upvoted 2 times

**SV1122** 3 years, 4 months ago

LOGIN_HISTORY returns login events within a specified time range. https://docs.snowflake.com/en/sql-reference/functions/login_history.html

upvoted 1 times

Question #413 *Topic 1*

What are common issues found by using the Query Profile? (Choose two.)

A. Identifying queries that will likely run very slowly before executing them

B. Locating queries that consume a high amount of credits

C. Identifying logical issues with the queries

D. Identifying inefficient micro-partition pruning

E. Data spilling to a local or remote disk

**Suggested Answer:** *D, E*

*Community vote distribution*

DE (100%)

**arpit_dataguy** Highly Voted 3 years, 4 months ago

DE :: https://docs.snowflake.com/en/user-guide/ui-query-profile.html#common-query-problems-identified-by-query-profile

upvoted 7 times

**Kate0204** Highly Voted 3 years, 5 months ago

I think DE

upvoted 6 times

**_yyukta** Most Recent 2 years, 2 months ago

DE are correct

upvoted 1 times

**e_l_i_a_s** 2 years, 3 months ago

If the local disk space is not sufficient, the spilled data is then saved to remote disks. Inefficient Pruning

upvoted 1 times

**miadiesel** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/ui-query-profile.html#statistics

upvoted 6 times

Question #414 *Topic 1*

The Snowflake Search Optimization Services supports improved performance of which kind of query?

A. Queries against large tables where frequent DML occurs

B. Queries against tables larger than 1 TB

C. Selective point lookup queries

D. Queries against a subset of columns in a table

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**_yyukta** Most Recent 2 years, 2 months ago

C. Selective point lookup queries

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

keyword: selective

upvoted 1 times

**MultiCloudIronMan** 3 years ago

Correct

upvoted 1 times

**SnowProCertDec22** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/search-optimization-service.html

upvoted 3 times

Question #415 *Topic 1*

Which file formats are supported for unloading data from Snowflake? (Choose two.)

A. Avro

B. JSON

C. ORC

D. XML

E. Delimited (CSV, TSV, etc.)

**Suggested Answer:** *B, E*

*Community vote distribution*

BE (100%)

**_yyukta** Most Recent 2 years, 2 months ago

BE are correct

upvoted 1 times

**e_l_i_a_s** 2 years, 3 months ago

Delimited files (CSV, TSV, etc.) Any valid delimiter is supported; default is comma (i.e. CSV). JSON Parquet https://docs.snowflake.com/en/user-guide/intro-summary-unloading#output-data-file-details

upvoted 2 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

JSON and delimited

upvoted 2 times

**miadiesel** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/intro-summary-unloading.html#output-data-file-details

upvoted 3 times

**Kate0204** 3 years, 5 months ago

CSV, JSON, Parquet

upvoted 4 times

Question #416 *Topic 1*

Which Snowflake tool would be BEST to troubleshoot network connectivity?

A. SnowCLI

B. SnowUI

C. SnowSQL

D. SnowCD

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**Kate0204** Highly Voted 3 years, 5 months ago

CD - Connectivity Diagnostics

upvoted 5 times

**_yyukta** Most Recent 2 years, 2 months ago

D. SnowCD is correct

upvoted 1 times

**e_l_i_a_s** 2 years, 3 months ago

SnowCD (i.e. Snowflake Connectivity Diagnostic Tool) helps users to diagnose and troubleshoot their network connection to Snowflake.

upvoted 1 times

**Sharu07** 2 years, 6 months ago

D - CORRECT ANSWER

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

D for sure. keyword: CD- connectivity diagnostics

upvoted 1 times

**MultiCloudIronMan** 3 years ago

Correct

upvoted 1 times

**miadiesel** 3 years, 4 months ago

SnowCD (i.e. Snowflake Connectivity Diagnostic Tool) helps users to diagnose and troubleshoot their network connection to Snowflake.

https://docs.snowflake.com/en/user-guide/snowcd.html#snowcd-connectivity-diagnostic-tool

upvoted 3 times

Question #417 *Topic 1*

Increasing the size of a virtual warehouse from an X-Small to an X-Large is an example of which of the following?

A. Right sizing

B. Concurrent sizing

C. Scaling out

D. Scaling up

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**_yyukta** Most Recent 2 years, 2 months ago

D. Scaling up

upvoted 2 times

**KnightVictor** 2 years, 11 months ago

Scaleup

upvoted 1 times

**MultiCloudIronMan** 3 years ago

Correct

upvoted 1 times

**SnowProCertDec22** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/warehouses-considerations.html#warehouse-resizing-improves-performance

upvoted 2 times

Question #418 *Topic 1*

What are ways to create and manage data shares in Snowflake? (Choose two.)

A. Through the Snowflake web interface (UI)

B. Through the DATA_SHARE=TRUE parameter

C. Through SQL commands

D. Through the ENABLE_SHARE=TRUE parameter

E. Using the CREATE SHARE AS SELECT * FROM TABLE command

**Suggested Answer:** *A, C*

*Community vote distribution*

AC (100%)

**_yyukta** Most Recent 2 years, 2 months ago

AC are correct

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**SV1122** 3 years, 4 months ago

Web Interface and SQL commands

upvoted 3 times

**AravindhTN** 3 years, 4 months ago

I think Answers correct

upvoted 2 times

Question #419 *Topic 1*

What is a characteristic of data micro-partitioning in Snowflake?

A. Micro-partitioning may introduce data skew

B. Micro-partitioning requires the definition of a partitioning schema

C. Micro-partitioning happens when the data is loaded

D. Micro-partitioning can be disabled within a Snowflake account

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**miadiesel** Highly Voted 3 years, 4 months ago

Micro-partitioning is automatically performed on all Snowflake tables. Tables are transparently partitioned using the ordering of the data as it is inserted/loaded. https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions.html#what-are-micro-partitions

upvoted 7 times

**PrasantSadangi** Most Recent 1 year, 7 months ago

C is the answer

upvoted 1 times

**_yyukta** 2 years, 2 months ago

C. Micro-partitioning happens when the data is loaded

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

C for sure

upvoted 1 times

**psilvabr** 3 years, 1 month ago

C is the right one.

upvoted 3 times

Question #420 *Topic 1*

Users with the ACCOUNTADMIN role can execute which of the following commands on existing users?

A. Can SHOW users DESCRIBE a given user, or ALTER or DROP a user

B. Can DEFINE users, DESCRIBE a given user, or ALTER or DELETE a user

C. Can SHOW users, INDEX a given user, or ALTER or DELETE a user

D. Can SHOW users, DEFINE a given user or ALTER, DROP, or MODIFY a user

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**_yyukta** Most Recent 2 years, 2 months ago

A. Can SHOW users DESCRIBE a given user, or ALTER or DROP a user

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**SV1122** 3 years, 4 months ago

Only these operations are allowed CREATE USER , ALTER USER , DROP USER , DESCRIBE USER https://docs.snowflake.com/en/sql-reference/sql/show-users.html

upvoted 4 times

**SnowProCertDec22** 3 years, 4 months ago

Selected Answer: A

by discard, you cannot use DELETE nor DEFINE nor MODIFY when you talk about users

upvoted 2 times

**SnowProCertDec22** 3 years, 4 months ago

by discard, you cannot use DELETE nor DEFINE nor MODIFY when you talk about usershttps://docs.snowflake.com/en/user-guide/admin-user-management.html

upvoted 2 times

**SnowProCertDec22** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/admin-user-management.html

upvoted 2 times

**halol** 3 years, 4 months ago

i think A

upvoted 1 times

Question #421 *Topic 1*

According to Snowflake best practice recommendations, which system-defined roles should be used to create custom roles? (Choose two.)

A. ACCOUNTADMIN

B. SYSADMIN

C. SECURITYADMIN

D. USERADMIN

E. ORGADMIN

**Suggested Answer:** *C, D*

*Community vote distribution*




**miadiesel** Highly Voted 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#system-defined-roles

upvoted 11 times

**d22770a** Most Recent 2 years ago

CD

upvoted 1 times

**joyboy23** 2 years, 2 months ago

Answer is CD Custom account roles can be created using the USERADMIN role (or a higher role) as well as by any role to which the CREATE ROLE privilege has been granted. https://docs.snowflake.com/en/user-guide/security-access-control-overview#label-access-control-overview-role-types Refer 'Custome Role' section in the above link

upvoted 1 times

**Lematthew31** 2 years, 6 months ago

AccountAdmin & SysAdmin

upvoted 2 times

**[Removed]** 3 years ago

Don't spam wrong answers - see the diagram and docs before votingCD is correct. Scroll above in the link where u see the diagram, and see the definition. You assign all roles to SYSADMIN but roles can be created by USERADMIN and SECURITYADMIN access only.

upvoted 4 times

**[Removed]** 3 years ago

Selected Answer: AB

Don't spam wrong answers - see the diagram and docs before voting

upvoted 4 times

**d22770a** 2 years ago

CD is correct. Scroll above in the link where u see the diagram, and see the definition. You assign all roles to SYSADMIN but roles can be created by USERADMIN and SECURITYADMIN access only.

upvoted 1 times

**Rajivnb** 3 years ago

SysAdmin : Does not have the power to create a Custom Role. A Custom Role can be assigned with SysAdmin Privilege but, it cannot create a Role. Security Admin and UserAdmin can create, Manage Users and Roles. Its not a good practice to use the AccountAdmin to create the Roles. Though it can be created, the role should be assigned to SysAdmin. https://www.chaosgenius.io/blog/snowflake-roles-access-control/#:~:text=SECURITYADMIN%20(Security%20Administrator)%3A&text=It%20has%20the%20ability%20to%20create%2C%20monitor%2C%20and%20manage%20users,via%20the%20system%20role%20hierarchy.

upvoted 1 times

**MultiCloudIronMan** 3 years, 4 months ago

https://www.bing.com/search?q=According+to+Snowflake+best+practice+recommendations%2C+which+system-defined+roles+should+be+used+to+create+custom+roles%3F&cvid=878ab2eb3e7a4c1eb414d5f94dfd4759&aqs=edge..69i57j69i11004.1528j0j9&FORM=ANAB01&PC=U531CD is correct answer

upvoted 1 times

**MultiCloudIronMan** 3 years, 4 months ago

Selected Answer: BD

https://www.bing.com/search?q=According+to+Snowflake+best+practice+recommendations%2C+which+system-defined+roles+should+be+used+to+create+custom+roles%3F&cvid=878ab2eb3e7a4c1eb414d5f94dfd4759&aqs=edge..69i57j69i11004.1528j0j9&FORM=ANAB01&PC=U531

upvoted 1 times

**MultiCloudIronMan** 3 years, 3 months ago

CD is correct answer

upvoted 1 times

**Dalo** 3 years, 5 months ago

Selected Answer: AB

show this picture https://docs.snowflake.com/ko/user-guide/security-access-control-overview#role-hierarchy-and-privilege-inheritance

upvoted 2 times

**Dalo** 3 years, 5 months ago

show this picture https://docs.snowflake.com/ko/user-guide/security-access-control-overview#role-hierarchy-and-privilege-inheritanceScroll above in the link you shared, and see the definition. You assign all roles to SYSADMIN but roles can be created by USERADMIN and SECURITYADMIN access only.

upvoted 2 times

**d22770a** 2 years ago

Scroll above in the link you shared, and see the definition. You assign all roles to SYSADMIN but roles can be created by USERADMIN and SECURITYADMIN access only.

upvoted 1 times

**MultiCloudIronMan** 3 years, 6 months ago

Answer

upvoted 1 times

**KarBiswa** 3 years, 7 months ago

Selected Answer: AB

https://docs.snowflake.com/en/user-guide/security-access-control-configure#label-security-custom-role

upvoted 4 times

**KarBiswa** 3 years, 7 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-configure#label-security-custom-roleSorry changing the option after reading the question multiple times got the flavour it is asking system defined role who can create custom roles in that case it is USERADMIN or higher so right option is SECURITYADMIN or USERADMINSYSADMIN can not define users or roles

upvoted 4 times

**KarBiswa** 3 years, 7 months ago

Sorry changing the option after reading the question multiple times got the flavour it is asking system defined role who can create custom roles in that case it is USERADMIN or higher so right option is SECURITYADMIN or USERADMIN

upvoted 2 times

**OTE** 3 years, 7 months ago

SYSADMIN can not define users or roles

upvoted 1 times

**KarBiswa** 3 years, 7 months ago

I would go for Accountadmin and SYSadmin as per https://docs.snowflake.com/en/user-guide/security-access-control-configure#label-security-custom-role

upvoted 2 times

**SnowProCertDec22** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#custom-rolesCustom account roles can be created using the USERADMIN role (or a higher role)

upvoted 3 times

**SnowProCertDec22** 3 years, 10 months ago

Selected Answer: CD

https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#custom-roles

upvoted 3 times

**SnowProCertDec22** 3 years, 10 months ago

Custom account roles can be created using the USERADMIN role (or a higher role)

upvoted 1 times

**Rob__C** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#system-defined-rolesQuestion 461 is also referring the same and i see the answer you provided there is SYSADMIN. Can you please confirmSYSADMIN role should be granted to custom roles after their creation. Custom account roles can be created using the USERADMIN role (or a higher role).

upvoted 1 times

**Rob__C** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#system-defined-roles

upvoted 1 times

**harshagc180** 3 years, 10 months ago

Question 461 is also referring the same and i see the answer you provided there is SYSADMIN. Can you please confirmSYSADMIN role should be granted to custom roles after their creation. Custom account roles can be created using the USERADMIN role (or a higher role).

upvoted 1 times

**sakis213** 3 years, 10 months ago

SYSADMIN role should be granted to custom roles after their creation. Custom account roles can be created using the USERADMIN role (or a higher role).

upvoted 3 times

**arpit_dataguy** 3 years, 10 months ago

Answer -- CD

upvoted 3 times

**arpit_dataguy** 3 years, 10 months ago

Answer -- CDQuestion 461 is also referring the same and i see the answer you provided there is SYSADMIN. Can you please confirm

upvoted 3 times

**harshagc180** 3 years, 10 months ago

Question 461 is also referring the same and i see the answer you provided there is SYSADMIN. Can you please confirm

upvoted 1 times

Question #422 *Topic 1*

What services are provided by the cloud services layer in Snowflake? (Choose two.)

A. Metadata management

B. Object authorization

C. Authentication

D. Query execution

E. Result caching

**Suggested Answer:** *A, C*

*Community vote distribution*

AC (100%)

**SnowProCertDec22** Highly Voted 2 years, 10 months ago

CLOUD SERVICES LAYER: • Authentication • Infrastructure management • Metadata management • Query parsing and optimization Access control

upvoted 5 times

**markpumc** Highly Voted 2 years, 8 months ago

A. Metadata management C. Authentication The cloud services layer in Snowflake provides several important services that are used by other parts of the system. Two of the main services provided by the cloud services layer are: A. Metadata management: The cloud services layer is responsible for managing the metadata that defines the structure and contents of databases and tables in Snowflake, including information about columns, data types, and indexes. C. Authentication: The cloud services layer is responsible for handling user authentication and authorization, ensuring that only authorized users have access to the data and metadata stored in Snowflake. Options B, D, and E refer to other parts of the Snowflake system and are not provided by the cloud services layer. Option B refers to object authorization, which is managed by the security layer in Snowflake. Option D refers to query execution, which is performed by the query processing layer. Option E refers to result caching, which can be performed by the query processing layer and the storage layer in Snowflake.

upvoted 5 times

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 5 months ago

Cloud layer is responsible for authentication and metadata mgmt

upvoted 1 times

**Rob__C** 2 years, 10 months ago

https://docs.snowflake.com/en/user-guide/intro-key-concepts.html#cloud-services

upvoted 2 times

Question #423 *Topic 1*

Which of the following commands are valid options for the VALIDATION_MODE parameter within the Snowflake COPY_INTO command? (Choose two.)

A. TRUE

B. RETURN_ERROR_SUM

C. RETURN_ALL_ERRORS

D. RETURN_[N]_ROWS

E. RETURN_FIRST__ERRORS

**Suggested Answer:** *C, D*

*Community vote distribution*

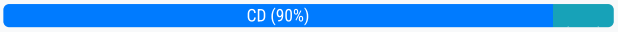


**SV1122** Highly Voted 2 years, 10 months ago

VALIDATION_MODE = RETURN_n_ROWS | RETURN_ERRORS | RETURN_ALL_ERRORS https://docs.snowflake.com/en/sql-reference/sql/copy-into-table.htmlThen there is an error on the question RETURN__ROWS is not a valid answer, shoud be like: RETURN_n_ROWS (e.g. RETURN_10_ROWS)RETURN__ROWS includes two "_" "_", I believe there's a hidden "n" in between. :)

upvoted 7 times

**SV1122** Highly Voted 2 years, 10 months ago

Selected Answer: CD

VALIDATION_MODE = RETURN_n_ROWS | RETURN_ERRORS | RETURN_ALL_ERRORS https://docs.snowflake.com/en/sql-reference/sql/copy-into-table.html

upvoted 7 times

**BigDataBB** 2 years, 9 months ago

Then there is an error on the question RETURN__ROWS is not a valid answer, shoud be like: RETURN_n_ROWS (e.g. RETURN_10_ROWS)

upvoted 3 times

**BigDataBB** 2 years, 9 months ago

Then there is an error on the question RETURN__ROWS is not a valid answer, shoud be like: RETURN_n_ROWS (e.g. RETURN_10_ROWS)RETURN__ROWS includes two "_" "_", I believe there's a hidden "n" in between. :)

upvoted 3 times

**ryotan** 2 years, 2 months ago

RETURN__ROWS includes two "_" "_", I believe there's a hidden "n" in between. :)

upvoted 2 times

**PrasantSadangi** Most Recent 1 year, 7 months ago

[ VALIDATION_MODE = RETURN_ _ROWS | RETURN_ERRORS | RETURN_ALL_ERRORS ]

upvoted 1 times

**Hopefan** 2 years, 5 months ago

CD there is no RETURN_ERROR_SUM

upvoted 1 times

**mike24061992** 2 years, 6 months ago

B. RETURN_ERROR_SUM C. RETURN_ALL_ERRORS Here is a brief explanation of these options: RETURN_ERROR_SUM: This option provides a summary of all validation errors encountered during the data load. It will not load any data into the table. RETURN_ALL_ERRORS: This option returns details of all validation errors encountered during the data load. It will not load any data into the table.

upvoted 1 times

**OTE** 2 years, 8 months ago

I believe this question is incorrect. Answer (C) is clear, but RETURN_ROWS (D) is not valid. Options are: VALIDATION_MODE = RETURN_ _ROWS | RETURN_ERRORS | RETURN_ALL_ERRORS

upvoted 1 times

**SnowProCertDec22** 2 years, 10 months ago

https://docs.snowflake.com/en/sql-reference/sql/copy-into-table.html#optional-parameters

upvoted 1 times

Question #424 *Topic 1*

Snowflake virtual warehouses are part of which layer of the Snowflake architecture?

A. Compute layer

B. Storage layer

C. Database layer

D. Cloud services layer

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 5 months ago

for sure compute layer

upvoted 1 times

**fahfouhi94** 2 years, 8 months ago

COMPUTE LAYER

upvoted 3 times

Question #425 *Topic 1*

Which of the following are characteristics of schemas used in Snowflake? (Choose two.)

A. A schema may contain one or more databases.

B. A database may contain one or more schemas.

C. A schema represents a logical grouping of database objects.

D. Each schema is contained within a virtual warehouse.

E. A table can span more than one schema.

**Suggested Answer:** *B, C*

*Community vote distribution*

BC (100%)

**Rob__C** Highly Voted 2 years, 4 months ago

B & C maybe ?

upvoted 10 times

**KnightVictor** Most Recent 1 year, 11 months ago

for sure B & C

upvoted 2 times

**Harpreetwal** 2 years ago

BC is correct

upvoted 2 times

**Harpreetwal** 2 years ago

B & C is correct

upvoted 2 times

**MultiCloudIronMan** 2 years ago

Correct

upvoted 2 times

**fahfouhi94** 2 years, 2 months ago

BC is the answer

upvoted 3 times

**AlexbDku** 2 years, 4 months ago

https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#securable-objects

upvoted 2 times

**AlexbDku** 2 years, 4 months ago

B&C of course https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#custom-roles

upvoted 2 times

Question #426 *Topic 1*

Which Snowflake objects can be used to reduce data storage costs for short-lived tables? (Choose two.)

A. Provisional tables

B. Temporary tables

C. Transient tables

D. Permanent tables

E. Lookup tables

**Suggested Answer:** *B, C*

*Community vote distribution*




**Harpreetwal** Most Recent 2 years ago

BC Correct

upvoted 1 times

**MultiCloudIronMan** 2 years ago

Correct

upvoted 1 times

**fahfouhi94** 2 years, 2 months ago

BC is correct

upvoted 3 times

Question #427 *Topic 1*

A user has unloaded data from Snowflake to a stage.

Which SQL command should be used to validate which data was loaded into the stage?

A. list @file_stage

B. show @file_stage

C. view @file_stage

D. verify @file_stage

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 1 times

**Harpreetwal** 2 years, 6 months ago

List command is correct

upvoted 1 times

**deadpool1108** 2 years, 7 months ago

A is correct.

upvoted 1 times

**okurok** 2 years, 7 months ago

https://docs.snowflake.com/en/sql-reference/sql/list

upvoted 2 times

**Jens_B** 2 years, 9 months ago

Correct

upvoted 1 times

Question #428 *Topic 1*

What are benefits of using the ACCESS_HISTORY view in the SNOWFLAKE database?

A. Identification of unused data

B. Identification of which roles have been used

C. Tracking of network policy usage

D. Highlighting of row access policy usage

E. Identification of who has read data

**Suggested Answer:** *A, E*

*Community vote distribution*




**Lematthew31** Most Recent 2 years ago

It could be also BDE

upvoted 1 times

**0e504b5** 2 years, 3 months ago

A E C - no, use this instead: https://docs.snowflake.com/en/sql-reference/account-usage/network_policies B - I would probably use QUERY_HISTORY for that instead https://docs.snowflake.com/en/sql-reference/account-usage/query_history D - I would use this instead: ROW_ACCESS_POLICIES https://docs.snowflake.com/en/sql-reference/account-usage/row_access_policies

upvoted 2 times

**baskar14** 2 years, 6 months ago

Specifies the DDL operation on a database, schema, table, view, and column. These operations also include statements that specify a row access policy on a table or view, a masking policy on a column, and tag updates (e.g. set a tag, change a tag value) on the object or column. USER_NAME column JSMITH

upvoted 1 times

**Marge23** 2 years, 7 months ago

[ { "objectDomain": "Stage", "objectId": 1, "objectName": "JRAMCELESTE", "stageKind": "User" } ]

upvoted 1 times

**shankarj68** 2 years, 9 months ago

A is the right answer. Check the benefit section: https://docs.snowflake.com/en/user-guide/access-history#benefits

upvoted 4 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 2 times

**Flava** 2 years, 10 months ago

Access History in Snowflake refers to when the user query reads column data and when the SQL statement performs a data write operation, such as INSERT, UPDATE, and DELETE along with variations of the COPY command, from the source data object to the target data object

upvoted 1 times

**Kelvin_Yu** 2 years, 11 months ago

ADE A: https://docs.snowflake.com/en/user-guide/access-history#label-access-history-benefits:~:text=Discover%20unused%20data%20to%20determine%20whether%20to%20archive%20or%20delete%20the%20data. D: https://docs.snowflake.com/en/user-guide/access-history#example-track-row-access-policy-references E: https://docs.snowflake.com/en/user-guide/access-history#examples-read-queries

upvoted 2 times

**bee5k** 2 years, 11 months ago

Only A. They did not ask for two.Dont Confuse

upvoted 1 times

**Harpreetwal** 3 years ago

AE looks correct

upvoted 2 times

**learner2023** 3 years ago

Benefits¶ Access history in Snowflake provides the following benefits pertaining to read and write operations: Data discovery Discover unused data to determine whether to archive or delete the data.

upvoted 2 times

**Kvk117** 3 years ago

Verified

upvoted 2 times

**GiselaS** 3 years, 1 month ago

E: "Access History in Snowflake refers to when the user query reads column data and when the SQL statement performs a data write operation" C: "These records facilitate regulatory compliance auditing and provide insights on popular and frequently accessed tables and columns"

upvoted 1 times

**okurok** 3 years, 1 month ago

**Selected Answer: DE**

D: https://docs.snowflake.com/en/sql-reference/account-usage/access_history#:~:text=Specifies%20information%20about%20the%20enforced%20masking%20policy%20set%20on%20the%20column%20and%20the%20enforced%20row%20access%20policy%20set%20on%20the%20table%2C%20including%20policies%20set%20on%20intermediate%20objects%20or%20columns.

E: https://docs.snowflake.com/en/sql-reference/account-usage/access_history#:~:text=The%20user%20who%20issued%20the%20query.

upvoted 1 times

**okurok** 3 years, 1 month ago

D: https://docs.snowflake.com/en/sql-reference/account-usage/access_history#:~:text=Specifies%20information%20about%20the%20enforced%20masking%20policy%20set%20on%20the%20column%20and%20the%20enforced%20row%20access%20policy%20set%20on%20the%20table%2C%20including%20policies%20set%20on%20intermediate%20objects%20or%20columns. E: https://docs.snowflake.com/en/sql-reference/account-usage/access_history#:~:text=The%20user%20who%20issued%20the%20query.also A https://docs.snowflake.com/en/user-guide/access-history#:~:text=Discover%20unused%20data

upvoted 1 times

**okurok** 3 years, 1 month ago

also A https://docs.snowflake.com/en/user-guide/access-history#:~:text=Discover%20unused%20data

upvoted 1 times

**KarBiswa** 3 years, 1 month ago

No mention of roles into the columns of the view into the user document

upvoted 1 times

**markpumc** 3 years, 2 months ago

The benefits of using the ACCESS_HISTORY view in the Snowflake database include: Identification of who has read data Tracking of which roles have been used Highlighting of row access policy usage By using the ACCESS_HISTORY view, organizations can get a better understanding of who has access to what data within their Snowflake environment, and ensure that data is being accessed and used in an appropriate manner.

upvoted 3 times

**sakis213** 3 years, 4 months ago

Access history in Snowflake provides the following benefits pertaining to read and write operations: Data discovery Discover unused data to determine whether to archive or delete the data.

upvoted 2 times

Question #429 *Topic 1*

Which of the following view types are available in Snowflake? (Choose two.)

A. Layered view

B. Secure view

C. External view

D. Embedded view

E. Materialized view

**Suggested Answer:** *B, E*

*Community vote distribution*

BE (100%)

**e_l_i_a_s** Most Recent 2 years, 3 months ago

https://docs.snowflake.com/en/user-guide/views-introduction

upvoted 1 times

**Harpreetwal** 3 years ago

BE Correct, Never heard any of the views like Layered, External or Embedded in any of the databases.

upvoted 1 times

**MultiCloudIronMan** 3 years, 2 months ago

Answer is correct

upvoted 2 times

**AravindhTN** 3 years, 4 months ago

Answer are correct

upvoted 2 times

Question #430 *Topic 1*

Which of the following statements describes a benefit of Snowflake's separation of compute and storage? (Choose two.)

A. Growth of storage and compute are tightly coupled.

B. Storage expands without the requirement to add more compute.

C. Compute can be scaled up or down without the requirement to add more storage.

D. Compute and storage can be scaled together.

E. Use of storage avoids disk spilling.

**Suggested Answer:** *B, C*

*Community vote distribution*

BC (100%)

**MultiCloudIronMan** Most Recent 2 years, 3 months ago

correct

upvoted 2 times

**Harpreetwal** 2 years, 6 months ago

Both can be managed saperately as per business needs. They are not bounded, this is what cloud is for

upvoted 2 times

**learner2023** 2 years, 6 months ago

For Snowflake, this is not an issue. You can grow and shrink the environment dynamically. The data storage grows and shrinks as you add or remove data, while the compute nodes can be ramped up or down, or turned off, as you require. You are not forced to pay for capacity up front, or kick other workloads off, or plan downtimes when ramping up your data warehouse capacity. That is the promise of cloud. This is why Snowflake is the real elastic Data Warehouse as a Service.

upvoted 2 times

**Hopefan** 2 years, 7 months ago

Need more help on this as D is confusing me! I guess B and C are correct, but D also seems ok ?

upvoted 1 times

Question #431 *Topic 1*

Which of the following languages can be used to implement Snowflake User Defined Functions (UDFs)? (Choose two.)

A. Ruby

B. Javascript

C. SQL

D. PERL

E. C#

**Suggested Answer:** *B, C*

*Community vote distribution*

BC (100%)

**Harpreetwal** Most Recent 2 years, 6 months ago

Correct

upvoted 2 times

**MultiCloudIronMan** 2 years, 6 months ago

Correct

upvoted 2 times

**SV1122** 2 years, 10 months ago

Java, Javascript, Python, SQL https://docs.snowflake.com/en/sql-reference/user-defined-functions.html

upvoted 2 times

Question #432 *Topic 1*

What is the default compression typo when unloading data from Snowflake?

A. Brotli

B. bzip2

C. Zstandard

D. gzip

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**SilviaS** Most Recent 10 months, 1 week ago

The file format also should be indicated. Snappy is not there, so, the only option is D

upvoted 1 times

**MultiCloudIronMan** 2 years ago

Correct

upvoted 1 times

**SV1122** 2 years, 4 months ago

By default, all unloaded data files are compressed using gzip, unless compression is explicitly disabled or one of the other supported compression methods is explicitly specified. https://docs.snowflake.com/en/user-guide/intro-summary-unloading.html#:~:text=By%20default%2C%20all%20unloaded%20data,compression%20methods%20is%20explicitly%20specified.

upvoted 2 times

Question #433 *Topic 1*

Which statement describes when a virtual warehouse can be resized?

A. A resize will affect running, queued, and new queries.

B. A resize can only be completed when the warehouse is in an auto-resume status.

C. A resize must be completed when the warehouse is suspended.

D. A resize can be completed at any time.

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**ramamama** Most Recent 2 years, 2 months ago

Ours i

upvoted 1 times

**MultiCloudIronMan** 3 years, 6 months ago

Correct

upvoted 1 times

**sakis213** 3 years, 10 months ago

A warehouse can be resized up or down at any time, including while it is running and processing statements.

upvoted 3 times

Question #434 *Topic 1*

What is the compressed size limit for semi-structured data loaded into a VARIANT data type using the COPY command?

A. 8 MB

B. 16 MB

C. 32 MB

D. 64 MB

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**Avi0407** Most Recent 1 year, 12 months ago

16MB is the answer

upvoted 1 times

**MultiCloudIronMan** 2 years, 6 months ago

Correct

upvoted 1 times

**pre_321** 2 years, 8 months ago

B is correct

upvoted 3 times

Question #435 *Topic 1*

User A cloned a schema and overwrote a schema that User B was working on. User B no longer has access to their version of the tables. However, this all occurred within the Time Travel retention period defined at the database level.

How should the missing tables be restored?

A. Use an UNDROP TABLE statement.

B. Use a CREATE TABLE AS SELECT statement

C. Rename the cloned schema and use an UNDROP SCHEMA statement.

D. Contact Snowflake Support to retrieve the data from Fail-safe

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**BobFar** Highly Voted 2 years, 6 months ago

C is correct! If an object with the same name already exists, UNDROP fails. Need to rename the existing object, which then enables us to restore the previous version of the object.

upvoted 5 times

**aws_god** Most Recent 2 years ago

https://docs.snowflake.com/en/sql-reference/sql/undrop-schema

upvoted 1 times

**MultiCloudIronMan** 3 years ago

Correct

upvoted 1 times

**pre_321** 3 years, 2 months ago

C is correct

upvoted 3 times

Question #436 *Topic 1*

How does Snowflake recommend handling the bulk loading of data batches from files already available in cloud storage?

A. Use Snowpipe.

B. Use the INSERT command.

C. Use an external table.

D. Use the COPY command.

**Suggested Answer:** *D*

*Community vote distribution*

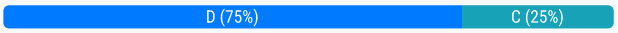


**nexerSnow** Most Recent 1 year, 11 months ago

answer should be C and D combined

upvoted 1 times

**BobFar** 2 years, 6 months ago

Answer D is correct. Step 1 create external stage Step 2 using copy into https://docs.snowflake.com/en/user-guide/data-load-gcs#:~:text=You%20can%20load%20directly%20from,manually%20or%20within%20a%20script.

upvoted 2 times

**MultiCloudIronMan** 2 years, 10 months ago

https://www.bing.com/search?q=How+does+Snowflake+recommend+handling+the+bulk+loading+of+data+batches+from+files+already+available+in+cloud+storage%3F&cvid=872b55f1708c4798b791492277cbcb09&aqs=edge..69i57j69i11004.2129j0j9&FORM=ANAB01&PC=U531

upvoted 1 times

**MultiCloudIronMan** 3 years ago

Correct

upvoted 2 times

**pre_321** 3 years, 2 months ago

Answer is correct

upvoted 3 times

Question #437 *Topic 1*

What is Snowflake's general guideline for files used to load data?

A. Files can be loaded directly into a table.

B. Any delimiter is supported: the default is a semicolon.

C. Electronic Data Interchange (EDI) is one of the supported semi-structured formats.

D. For delimited files, the default character set is UTF-8.

**Suggested Answer:** *D*

*Community vote distribution*

D (96%) (4

**SnowProCertDec22** Highly Voted 3 years, 4 months ago

Selected Answer: D

by discard, not A because you have to load the data on a stage before being able to load it into a table. not b because comma is the default delimiter. not c because edi is not supported.

upvoted 18 times

**SnowProCertDec22** Highly Voted 3 years, 4 months ago

by discard, not A because you have to load the data on a stage before being able to load it into a table. not b because comma is the default delimiter. not c because edi is not supported.Great explanation.Files are first copied ("staged") to an internal (Snowflake) stage, then loaded into a table.

upvoted 18 times

**AravindhTN** 3 years, 4 months ago

Great explanation.

upvoted 1 times

**SnowProCertDec22** 3 years, 4 months ago

Files are first copied ("staged") to an internal (Snowflake) stage, then loaded into a table.

upvoted 1 times

**ShawnZh** Most Recent 1 year, 1 month ago

D is best fit.

upvoted 1 times

**AlphaCentauriBear** 2 years, 2 months ago

A is also correct. https://docs.snowflake.com/en/sql-reference/sql/copy-into-table COPY INTO command allows to copy files directly from remote location bypassing the stage. COPY INTO [ .] FROM { internalStage | externalStage | externalLocation }Well, yes, but it is not recommended.

upvoted 1 times

**AlphaCentauriBear** 2 years, 2 months ago

Selected Answer: A

A is also correct.

https://docs.snowflake.com/en/sql-reference/sql/copy-into-table COPY INTO command allows to copy files directly from remote location bypassing the stage.

COPY INTO [ .] FROM { internalStage | externalStage | externalLocation }

upvoted 1 times

**yaho5** 2 years, 1 month ago

Well, yes, but it is not recommended.

upvoted 1 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 1 times

**KnightVictor** 2 years, 11 months ago

for sure D

upvoted 1 times

**snowcert** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/intro-summary-loading.html

upvoted 3 times

**Rob__C** 3 years, 4 months ago

https://docs.snowflake.com/en/user-guide/intro-summary-loading.html

upvoted 1 times

Question #438 *Topic 1*

How does a Snowflake user execute an anonymous block of code?

A. The user must run the CALL command to execute the block.

B. The statements that define the block must also execute the block.

C. The SUBMIT command must run immediately after the block is defined

D. The block must be saved to a worksheet and executed using a connector.

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**SnowProCertDec22** Highly Voted 3 years, 4 months ago

The BEGIN … END statement that defines the block also executes the block. (You don't run a separate CALL command to execute the block.) https://docs.snowflake.com/en/developer-guide/snowflake-scripting/blocks.html#using-an-anonymous-block

upvoted 5 times

**e_l_i_a_s** Most Recent 2 years, 3 months ago

If you don't want to store the block in a stored procedure in the database, you can define and use an anonymous block. An anonymous block is a block that is not part of a stored procedure. You define the block as a separate, standalone SQL statement. The BEGIN statement that defines the block also executes the block. (You don't run a separate CALL command to execute the block.) https://docs.snowflake.com/en/developer-guide/snowflake-scripting/blocks#using-an-anonymous-block

upvoted 3 times

**Rob__C** 3 years, 4 months ago

https://docs.snowflake.com/en/developer-guide/snowflake-scripting/blocks.html#using-an-anonymous-block

upvoted 3 times

Question #439 *Topic 1*

When unloading data from Snowflake, the user executes a COPY INTO [location] command into an internal stage.

What additional command is required to load the file onto the local file system?

A. GET

B. LIST

C. PUT

D. REMOVE

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**EmiB** Highly Voted 2 years, 1 month ago

A - GET (stage to file system) C - PUT (file system to stage)

upvoted 5 times

**MultiCloudIronMan** Most Recent 2 years ago

Verified

upvoted 1 times

**miadiesel** 2 years, 3 months ago

Downloads data files from one of the following Snowflake stages to a local directory/folder on a client machine https://docs.snowflake.com/en/sql-reference/sql/get.html#get

upvoted 3 times

**Rob__C** 2 years, 4 months ago

https://docs.snowflake.com/en/user-guide/data-unload-overview.html#bulk-unloading-process

upvoted 3 times

**yogen** 2 years, 4 months ago

A - GET

upvoted 3 times

Question #440 *Topic 1*

A Snowflake user has a query that is running for a long time. When the user opens the query profiler, it indicates that a lot of data is spilling to disk.

What is causing this to happen?

A. The result cache is almost full and is unable to hold the results.

B. The cloud storage staging area is not sufficient to hold the data results.

C. Clustering has not been applied to the table so the table is not optimized.

D. The warehouse memory is not sufficient to hold the intermediate query results.

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**sakis213** Highly Voted 2 years, 10 months ago

D is the correct answer

upvoted 5 times

**BobFar** Most Recent 2 years ago

Correct answer is D. The warehouse memory is not sufficient to hold the intermediate query results. For some operations (e.g. duplicate elimination for a huge data set), the amount of memory available for the compute resources used to execute the operation might not be sufficient to hold intermediate results. As a result, the query processing engine will start spilling the data to local disk. If the local disk space is not sufficient, the spilled data is then saved to remote disks. https://docs.snowflake.com/en/user-guide/ui-query-profile#queries-too-large-to-fit-in-memory Local disk space is located on the compute nodes that are used to execute Snowflake queries and remote disk is the permanent storage location for all Snowflake data.

upvoted 1 times

**MultiCloudIronMan** 2 years, 6 months ago

Verified

upvoted 1 times

**Rob__C** 2 years, 10 months ago

https://community.snowflake.com/s/article/Performance-impact-from-local-and-remote-disk-spilling

upvoted 2 times

Question #441 *Topic 1*

What is the MOST performant file format for loading data in Snowflake?

A. CSV (Unzipped)

B. Parquet

C. CSV (Gzipped)

D. ORC

**Suggested Answer:** *C*

*Community vote distribution*

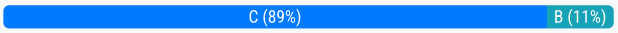


**leozhang** Highly Voted 2 years, 3 months ago

Loading from Gzipped CSV is several times faster than loading from ORC and Parquet at an impressive 15 TB/Hour. While 5-6 TB/hour is decent if your data is originally in ORC or Parquet, don't go out of your way to CREATE ORC or Parquet files from CSV in the hope that it will load Snowflake faster. Loading data into fully structured (columnarized) schema is ~10-20% faster than landing it into a VARIANT.
https://community.snowflake.com/s/article/How-to-Load-Terabytes-Into-Snowflake-Speeds-Feeds-and-TechniquesHere is the reference:
https://www.snowflake.com/en/engineering-blog/loading-terabytes-into-snowflake-speeds-feeds-techniques/

upvoted 7 times

**leozhang** Highly Voted 2 years, 3 months ago

Selected Answer: C

Loading from Gzipped CSV is several times faster than loading from ORC and Parquet at an impressive 15 TB/Hour. While 5-6 TB/hour is decent if your data is originally in ORC or Parquet, don't go out of your way to CREATE ORC or Parquet files from CSV in the hope that it will load Snowflake faster. Loading data into fully structured (columnarized) schema is ~10-20% faster than landing it into a VARIANT.

https://community.snowflake.com/s/article/How-to-Load-Terabytes-Into-Snowflake-Speeds-Feeds-and-Techniques

upvoted 7 times

**ShawnZh** Most Recent 1 year, 1 month ago

Here is the reference: https://www.snowflake.com/en/engineering-blog/loading-terabytes-into-snowflake-speeds-feeds-techniques/

upvoted 1 times

**ShawnZh** 1 year, 1 month ago

With recent advancements, the most performant file format for loading data into Snowflake is Parquet, when used with Snowflake's vectorized scanner. This is a shift from the past, where compressed CSV was often the fastest. Performance Comparison Recent benchmarks from Snowflake show that the vectorized scanner for Parquet significantly outperforms other formats in both speed and cost-efficiency. Parquet (Vectorized Scanner): This is now the fastest method. In a benchmark test, it took 626 seconds to load a large dataset. CSV: Previously considered the top performer, a compressed CSV file took 1,063 seconds to load the same dataset. JSON and Standard Parquet: These formats were considerably slower in the same test, taking over 2,000 seconds.

upvoted 1 times

**MultiCloudIronMan** 2 years ago

Selected Answer: C

Verified

upvoted 1 times

**MultiCloudIronMan** 2 years ago

Verifiedshould be B While CSV (Gzipped) (option C) is a commonly used format and can be efficient in terms of storage space due to compression, it is not as performant as Parquet (option B) for loading data in Snowflake.

upvoted 1 times

**MultiCloudIronMan** 1 year, 10 months ago

should be B While CSV (Gzipped) (option C) is a commonly used format and can be efficient in terms of storage space due to compression, it is not as performant as Parquet (option B) for loading data in Snowflake.

upvoted 1 times

Question #442 *Topic 1*

Which chart type does Snowsight support to visualize worksheet data?

A. Box plot

B. Bubble chart

C. Pie chart

D. Scatterplot

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**SV1122** Highly Voted 3 years, 10 months ago

Snowsight supports the following types of charts: Bar charts Line charts Scatterplots Heat grids Scorecards https://docs.snowflake.com/en/user-guide/ui-snowsight-visualizations.html

upvoted 7 times

**Rameez1** Most Recent 2 years, 4 months ago

Snowsight supports the following types of charts: Bar charts Line charts Scatterplots Heat grids Scorecards

upvoted 1 times

**MultiCloudIronMan** 3 years, 6 months ago

Correct

upvoted 1 times

**SnowProCertDec22** 3 years, 10 months ago

https://docs.snowflake.com/en/user-guide/ui-snowsight-visualizations.html#:~:text=Snowsight%20supports%20the%20following%20types,Scatterplots

upvoted 3 times

Question #443 *Topic 1*

Which result shows efficient pruning?

A. Partitions scanned is greater than partitions total.

B. Partitions scanned is less than partitions total.

C. Partitions scanned is equal to the partitions total.

D. Partitions scanned is greater than or equal to the partitions total.

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**AlexbDku** Highly Voted 3 years, 4 months ago

B https://docs.snowflake.com/en/user-guide/ui-query-profile.html#inefficient-pruning

upvoted 9 times

**gizzamo** Most Recent 2 years, 4 months ago

https://docs.snowflake.com/en/user-guide/ui-query-profile#inefficient-pruning ... The efficiency of pruning can be observed by comparing Partitions scanned and Partitions total statistics in the TableScan operators. If the former is a small fraction of the latter, pruning is efficient. If not, the pruning did not have an effect. ...

upvoted 2 times

**MultiCloudIronMan** 2 years, 9 months ago

correct

upvoted 2 times

**KnightVictor** 2 years, 11 months ago

for sure B. thats the primary purpose of pruning... to scan less partitions

upvoted 2 times

**MultiCloudIronMan** 3 years ago

B is correct

upvoted 2 times

**EmiB** 3 years, 1 month ago

B - correct

upvoted 2 times

**AravindhTN** 3 years, 4 months ago

Answer Should be B (scanned partitions should be less than total partitions)

upvoted 4 times

**sachchhab** 3 years, 4 months ago

B-https://docs.snowflake.com/en/user-guide/ui-query-profile.html#inefficient-pruning

upvoted 4 times