EXAMTOPICS

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



CERTIFICATION TEST

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

Question #1 *Topic 1*

What built-in Snowflake features make use of the change tracking metadata for a table? (Choose two.)

A. The MERGE command

B. The UPSERT command

C. The CHANGES clause

D. A STREAM object

E. Thee CHANGE_DATA_CAPTURE command

Suggested Answer: *C, D*

Community vote distribution

CD (100%)

JRayan Highly Voted 3 years, 6 months ago

The built-in Snowflake features that make use of the change tracking metadata for a table are the CHANGES clause and a STREAM object. Therefore, options (C) and (D) are the correct answers.

upvoted 10 times

NachoPrendes Most Recent 2 years, 2 months ago

CD

upvoted 1 times

cui_li 2 years, 4 months ago

C and D: Currently, the following must be true before change tracking metadata is recorded: Tables. Either enable change tracking on the table (using ALTER TABLE … CHANGE_TRACKING = TRUE), or create a stream on the table (using CREATE STREAM)

upvoted 2 times

Vipin_Singham 2 years, 6 months ago

Is it A and D. Merge command is used to upsert and delete based on Metadata informations from stream objects.So I think,it's A and D...

upvoted 1 times

Question #2 *Topic 1*

When using the Snowflake Connector for Kafka, what data formats are supported for the messages? (Choose two.)

A. CSV

B. XML

C. Avro

D. JSON

E. Parquet

Suggested Answer: *C, D*

Community vote distribution

CD (100%)

NachoPrendes Most Recent 2 years, 2 months ago

https://docs.snowflake.com/en/user-guide/kafka-connector-overview#:~:text=Each%20Kafka%20message%20is%20passed%20to%20Snowflake%20in%20JSON%20format%20or%20Avro%20format

upvoted 1 times

krishnak2244 3 years, 3 months ago

Ans: C,D Explanation: Each Kafka message is passed to Snowflake in JSON format or Avro format. The Kafka connector stores that formatted information in a single column of type VARIANT. The data is not parsed, and the data is not split into multiple columns in the Snowflake table. Link: https://docs.snowflake.com/en/user-guide/kafka-connector-overview

upvoted 2 times

Jay_98_11 3 years, 5 months ago

correct

upvoted 1 times

JRayan 3 years, 6 months ago

The Snowflake Connector for Kafka supports the JSON and Avro data formats for the messages. Therefore, options (C) and (D) are the correct answers. Snowflake does not support CSV, XML, or Parquet data formats for Kafka messages

upvoted 4 times

Question #3 *Topic 1*

At which object type level can the APPLY MASKING POLICY, APPLY ROW ACCESS POLICY and APPLY SESSION POLICY privileges be granted?

A. Global

B. Database

C. Schema

D. Table

Suggested Answer: *A*

Community vote distribution

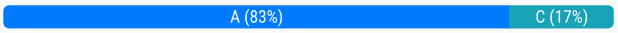

serg_khar Highly Voted 3 years, 6 months ago

my choice is A: https://docs.snowflake.com/en/user-guide/security-access-control-privileges#global-privileges

upvoted 9 times

Balakt Most Recent 1 year, 8 months ago

Schema

upvoted 1 times

MSIDDIQUI18 1 year, 11 months ago

Selected Answer: A APPLY MASKING POLICY Global Grants the ability to set a Column-level Security masking policy on a table or view column and to set a masking policy on a tag. This global privilege also allows executing the DESCRIBE operation on tables and views. APPLY ROW ACCESS POLICY Global Grants the ability to add and drop a row access policy on a table or view. This global privilege also allows executing the DESCRIBE operation on tables and views. APPLY SESSION POLICY Global Grants the ability to set or unset a session policy on an account or user.

upvoted 1 times

NachoPrendes 2 years, 2 months ago

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

upvoted 1 times

Gwo 2 years, 7 months ago

A is correct. Global object level. https://docs.snowflake.com/en/user-guide/security-access-control-privileges

upvoted 3 times

BDaly 2 years, 11 months ago

Table is correct, the listed privileges AAPLY to tables and views. While it is a global permission, the OBJECT they are applied to are tables/views. Hence D

upvoted 1 times

Arumugham 3 years, 1 month ago

Answer is A

upvoted 1 times

hillcat111 3 years, 3 months ago

Answer is A and is validated

upvoted 1 times

omsingh 3 years, 3 months ago

schema is correct answerit should be A, global

upvoted 1 times

omsingh 3 years, 3 months ago

schema is correct answer

upvoted 1 times

omsingh 3 years, 3 months ago

it should be A, global

upvoted 1 times

JRayan 3 years, 3 months ago

Global is privilege not object. Hence, C. Schema(object) is the right option.

upvoted 4 times

JRayan 3 years, 3 months ago

Global is privilege not object. Hence, C. Schema(object) is the right option.Agree, the way question is composed, I would say C.

upvoted 4 times

Corleone433 3 years, 1 month ago

Agree, the way question is composed, I would say C.

upvoted 1 times

henrik0 3 years, 6 months ago

seems as both global and schema are valid options

upvoted 1 times

qqcvd 3 years, 6 months ago

These are global privileges to grant, see : https://docs.snowflake.com/en/sql-reference/sql/grant-privilege

upvoted 4 times

Question #4 *Topic 1*

An Architect uses COPY INTO with the ON_ERROR=SKIP_FILE option to bulk load CSV files into a table called TABLEA, using its table stage. One file named file5.csv fails to load. The Architect fixes the file and re-loads it to the stage with the exact same file name it had previously.
Which commands should the Architect use to load only file5.csv file from the stage? (Choose two.)

A. COPY INTO tablea FROM @%tablea RETURN_FAILED_ONLY = TRUE;

B. COPY INTO tablea FROM @%tablea;

C. COPY INTO tablea FROM @%tablea FILES = ('file5.csv');

D. COPY INTO tablea FROM @%tablea FORCE = TRUE;

E. COPY INTO tablea FROM @%tablea NEW_FILES_ONLY = TRUE;

F. COPY INTO tablea FROM @%tablea MERGE = TRUE;

Suggested Answer: *B, C*

Community vote distribution

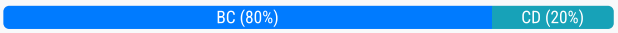

serg_khar Highly Voted 3 years, 6 months ago

Key point here - "ON_ERROR=SKIP_FILE". It means that file hasn't loaded. I we want to load the file we have to point it explicitly - this is C option or load all unloaded files (with file5.csv); its B option. If we use FORCE we will load again ALL data from the stage - it doesn't look as a good option.

upvoted 6 times

rajivgupta780 Most Recent 10 months, 3 weeks ago

A - Is invalid B - Will skill all uploaded file C - Will force to reload all file. D- Will make sure only File5.csv gets reloaded. E - Invalid F - Invalid.

upvoted 1 times

Gandharv75 1 year, 10 months ago

Option B and C are correct. Question says bulk load files, so there may be files other than fil5.csv. When File5.csv was modified and re-uploaded a new checksum was generated by snowflake, so B. COPY INTO tablea FROM @%tablea- will work as it won't load the files already loaded first time except file5.csv which has a new checksum now. Also, C. COPY INTO tablea FROM @%tablea FILES = ('file5.csv'); works because it specifically requests file5.csv to be loaded only(which already has anew checksum,so will be loaded successfully) Hence B and C only..

upvoted 1 times

ggorzki 2 years, 3 months ago

B and C Parameter NEW_FILES_ONLY does not exist. File is already changed in the stage, so FORCE = TRUE is not needed. Besides, it would load duplicates from the other files that reside in the stage.

upvoted 1 times

laksnarn 2 years, 3 months ago

C & D Here's the reasoning: C. COPY INTO tablea FROM @%tablea FILES = ('file5.csv'); This command specifically targets the file file5.csv, ensuring that only this file is loaded. D. COPY INTO tablea FROM @%tablea FORCE = TRUE; This command forces the reloading of the file, ignoring any previous load history and thus ensuring the fixed file5.csv is loaded even if the system might otherwise think it has already been processed.No D is not correct , it will load all the files again ! due to snowflake document, changing the files content, generated new checksum , so no need to force it https://docs.snowflake.com/en/sql-reference/sql/copy-into-table#reloading-files:~:text=generates%20a%20new-,checksum,-.

upvoted 3 times

laksnarn 2 years, 3 months ago

C & D Here's the reasoning:

C. COPY INTO tablea FROM @%tablea FILES = ('file5.csv'); This command specifically targets the file file5.csv, ensuring that only this file is loaded.

D. COPY INTO tablea FROM @%tablea FORCE = TRUE; This command forces the reloading of the file, ignoring any previous load history and thus ensuring the fixed file5.csv is loaded even if the system might otherwise think it has already been processed.

upvoted 3 times

BobFar 1 year, 1 month ago

No D is not correct , it will load all the files again ! due to snowflake document, changing the files content, generated new checksum , so no need to force it https://docs.snowflake.com/en/sql-reference/sql/copy-into-table#reloading-files:~:text=generates%20a%20new-,checksum,-.

upvoted 1 times

ivanjs44 2 years, 5 months ago

The key of this question is that say: The Architect fixes the file and re-loads Here in Snowflake documentation: https://docs.snowflake.com/en/sql-reference/sql/copy-into-table#usage-notes To reload the data, you must either specify FORCE = TRUE or ******modify the file and stage it again, which generates a new checksum.****

upvoted 1 times

cornz55 2 years, 10 months ago

The answer is C and D. You cannot choose C and take any other answer other than D because the file was previously loaded (with failure). So C will never work again unless there is D.

upvoted 1 times

jjordan 2 years, 10 months ago

there is not NEW_FILES_ONLY options.

upvoted 3 times

anikl 2 years, 11 months ago

There is no NEW_FILES_ONLY option. B and C are the correct

upvoted 3 times

jabibo 3 years, 1 month ago

B and C

upvoted 1 times

chagantik90 3 years, 1 month ago

+1 serg_khar, should be B and C

upvoted 1 times

Arbind_Chandra 3 years, 3 months ago

B and C are the correct answers.

upvoted 3 times

Needium 3 years, 3 months ago

The correct answer is B&C

upvoted 2 times

SrideviJ 3 years, 4 months ago

BC are the correct answers I think

upvoted 2 times

henrik0 3 years, 6 months ago

I think purge and force are the currect answers

upvoted 1 times

Question #5 *Topic 1*

A large manufacturing company runs a dozen individual Snowflake accounts across its business divisions. The company wants to increase the level of data sharing to support supply chain optimizations and increase its purchasing leverage with multiple vendors.
The company's Snowflake Architects need to design a solution that would allow the business divisions to decide what to share, while minimizing the level of effort spent on configuration and management. Most of the company divisions use Snowflake accounts in the same cloud deployments with a few exceptions for European-based divisions.
According to Snowflake recommended best practice, how should these requirements be met?

A. Migrate the European accounts in the global region and manage shares in a connected graph architecture. Deploy a Data Exchange.

B. Deploy a Private Data Exchange in combination with data shares for the European accounts.

C. Deploy to the Snowflake Marketplace making sure that invoker_share() is used in all secure views.

D. Deploy a Private Data Exchange and use replication to allow European data shares in the Exchange.

Suggested Answer: *D*

Community vote distribution

D (100%)

MultiCloudGuru Most Recent 1 year, 4 months ago
Answer is D
upvoted 1 times

prshntdxt7 2 years, 1 month ago
answer is D
upvoted 1 times

hillcat111 2 years, 9 months ago
Answer is D and is validated
upvoted 3 times

qqcvd 2 years, 9 months ago
correct
upvoted 1 times

hillcat111 2 years, 10 months ago
i think its BYou can not share the data into different region without replication. so option D seems correct.
upvoted 1 times

hillcat111 2 years, 10 months ago
i think its B
upvoted 1 times

Arbind_Chandra 2 years, 9 months ago
You can not share the data into different region without replication. so option D seems correct.
upvoted 3 times

Jay_98_11 2 years, 11 months ago
correct
upvoted 2 times

Question #6 *Topic 1*

A user has the appropriate privilege to see unmasked data in a column.
If the user loads this column data into another column that does not have a masking policy, what will occur?

A. Unmasked data will be loaded in the new column.

B. Masked data will be loaded into the new column.

C. Unmasked data will be loaded into the new column but only users with the appropriate privileges will be able to see the unmasked data.

D. Unmasked data will be loaded into the new column and no users will be able to see the unmasked data.

Suggested Answer: *A*

Community vote distribution

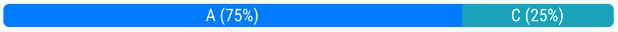

078f5e0 Most Recent 1 year, 8 months ago

Option A: Here is the explanation: Use caution when inserting values from a source column that has a masking policy on the source column to a target column without a masking policy on the target column. Since a masking policy is set on the source column, a role that views unmasked column data can insert unmasked data into another column, where any role with sufficient privileges on the table or view can see the value.

upvoted 1 times

laksnarn 2 years, 3 months ago

A. Unmasked data will be loaded in the new column. Here's why: When a user with appropriate privileges to see unmasked data loads data from a column with a masking policy into another column without a masking policy, the data loaded into the new column will be unmasked. This is because the user's privileges allow them to access the original data without the masking applied. Since the new column does not have a masking policy, it will simply store the data as it was accessed by the user, which in this case is unmasked. Therefore, the new column will contain the unmasked data regardless of the masking policy on the original column, and no additional masking will be applied to the new column unless explicitly defined.

upvoted 1 times

DBSS 2 years, 5 months ago

C Use caution when inserting values from a source column that has a masking policy on the source column to a target column without a masking policy on the target column. Since a masking policy is set on the source column, a role that views unmasked column data can insert unmasked data into another column, where any role with sufficient privileges on the table or view can see the value. If a role that sees masked data in the source column inserts those values into a target column, the inserted values remain masked. https://docs.snowflake.com/en/user-guide/security-column-intro

upvoted 2 times

jjordan 2 years, 9 months ago

Definitely A. "Since a masking policy is set on the source column, a role that views unmasked column data can insert unmasked data into another column, where any role with sufficient privileges on the table or view can see the value.". Doc: https://docs.snowflake.com/en/user-guide/security-column-intro

upvoted 1 times

papi97 2 years, 10 months ago

Answer is C

upvoted 1 times

chagantik90 3 years, 1 month ago

A is correct

upvoted 2 times

hillcat111 3 years, 3 months ago

Answer is A and is validated

upvoted 1 times

Needium 3 years, 4 months ago

Definitely A. "Since a masking policy is set on the source column, a role that views unmasked column data can insert unmasked data into another column, where any role with sufficient privileges on the table or view can see the value.". Doc: https://docs.snowflake.com/en/user-guide/security-column-intro

upvoted 4 times

Alachramkowa 3 years, 5 months ago

I copied masked vaues to anothe columns w/o masking policy. Vote for "A" option

upvoted 2 times

Jay_98_11 3 years, 5 months ago

"Since a masking policy is set on the source column, a role that views unmasked column data can insert unmasked data into another column, where any role with sufficient privileges on the table or view can see the value.". Doc: https://docs.snowflake.com/en/user-guide/security-column-introYour explanation only confirms A as the answerYour explanation point to A as the correct answer.

upvoted 1 times

Jay_98_11 3 years, 5 months ago

Selected Answer: C

"Since a masking policy is set on the source column, a role that views unmasked column data can insert unmasked data into another column, where any role with sufficient privileges on the table or view can see the value.". Doc: https://docs.snowflake.com/en/user-guide/security-column-intro

upvoted 1 times

Needium 3 years, 4 months ago

Your explanation only confirms A as the answer

upvoted 2 times

jjordan 2 years, 9 months ago

Your explanation point to A as the correct answer.

upvoted 1 times

Question #7 *Topic 1*

How can an Architect enable optimal clustering to enhance performance for different access paths on a given table?

A. Create multiple clustering keys for a table.

B. Create multiple materialized views with different cluster keys.

C. Create super projections that will automatically create clustering.

D. Create a clustering key that contains all columns used in the access paths.

Suggested Answer: *B*

Community vote distribution

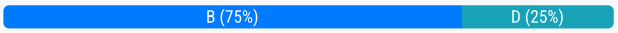

Indraneil2 Highly Voted 2 years, 4 months ago

Creating the materialized view with Snowflake allows you to specify the new clustering key, which enables Snowflake to reorganize the data during the initial creation of the materialized view.

upvoted 5 times

Yogendra_examtopic Most Recent 1 year, 8 months ago

Creating multiple materialized views with different cluster keys allows you to optimize for various access paths by distributing the data in ways that best suit each specific query pattern. This approach can significantly improve query performance as each materialized view can be tailored to different use cases, providing efficient access and faster retrieval times for specific queries.

upvoted 3 times

Balakt 1 year, 8 months ago

D is correct

upvoted 1 times

Makabaka 1 year, 11 months ago

B THE PROBLEM: MULTIPLE ACCESS PATHS THE SOLUTION: CLUSTERED MATERIALIZED VIEW

upvoted 1 times

hillcat111 2 years, 3 months ago

Answer is B and is validated

upvoted 3 times

hillcat111 2 years, 4 months ago

i think its D

upvoted 1 times

Question #8 *Topic 1*

Company A would like to share data in Snowflake with Company B. Company B is not on the same cloud platform as Company A.
What is required to allow data sharing between these two companies?

A. Create a pipeline to write shared data to a cloud storage location in the target cloud provider.

B. Ensure that all views are persisted, as views cannot be shared across cloud platforms.

C. Setup data replication to the region and cloud platform where the consumer resides.

D. Company A and Company B must agree to use a single cloud platform: Data sharing is only possible if the companies share the same cloud provider.

Suggested Answer: *C*

Community vote distribution

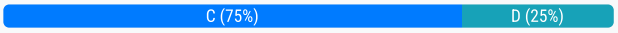

uncledana Most Recent 8 months, 1 week ago
C - data replication is the only one.
upvoted 1 times

caio_tertuliano 1 year, 7 months ago
Despite saying it comes from different organizations, the most fitted option would be C
upvoted 2 times

nareeshk 1 year, 8 months ago
All options have issues.. The correct answer would be to use listings with auto fulfilment which would take care of replication. Option C says replication on its own which implies direct database replication, and is not allowed between organisations... Since we are interested in Snowflake native offering, I have chosen option D as the potential solution, but it should have been option if its choice is correctly worded
upvoted 1 times

cloud_ranger 1 year, 10 months ago
Answer A. It cannot be answer C, because the companies would need to be in the "within the same organization", according to Snowflake documentation for replication.I don't think it means a direct replication to B's account. A will create an account in the target region and cloud platform for the replication, then share it with B.
upvoted 2 times

cloud_ranger 1 year, 10 months ago
Answer A. It cannot be answer C, because the companies would need to be in the "within the same organization", according to Snowflake documentation for replication.
upvoted 2 times

ywan1600 1 year, 6 months ago
I don't think it means a direct replication to B's account. A will create an account in the target region and cloud platform for the replication, then share it with B.
upvoted 2 times

hillcat111 2 years, 9 months ago
Answer is C and is validated
upvoted 3 times

SrideviJ 2 years, 10 months ago
Option C is correct since Cross cloud and cross platform data sharing is possible with replication
upvoted 1 times

Question #9 *Topic 1*

What are some of the characteristics of result set caches? (Choose three.)

A. Time Travel queries can be executed against the result set cache.

B. Snowflake persists the data results for 24 hours.

C. Each time persisted results for a query are used, a 24-hour retention period is reset.

D. The data stored in the result cache will contribute to storage costs.

E. The retention period can be reset for a maximum of 31 days.

F. The result set cache is not shared between warehouses.

Suggested Answer: *B, C, E*

Community vote distribution

BCE (100%)

uncledana Most Recent 8 months, 1 week ago

https://docs.snowflake.com/en/user-guide/querying-persisted-results

upvoted 1 times

BobFar 1 year, 1 month ago

BCE are correct Each time the persisted result for a query is reused, Snowflake resets the 24-hour retention period for the result, up to a maximum of 31 days from the date and time that the query was first executed. After 31 days, the result is purged and the next time the query is submitted, a new result is generated and persisted. https://docs.snowflake.com/en/user-guide/querying-persisted-results

upvoted 1 times

MSIDDIQUI18 1 year, 11 months ago

BCE are correct due to the following reasons: For persisted query results of all sizes, the cache expires after 24 hours. Each time the persisted result for a query is reused, Snowflake resets the 24-hour retention period for the result, up to a maximum of 31 days from the date and time that the query was first executed. After 31 days, the result is purged and the next time the query is submitted, a new result is generated and persisted.

upvoted 1 times

67422cb 2 years, 4 months ago

How is F not a correct answer? It is more correct than B, since a result cache be kicked out of memory before 24 hours if there are newer results coming in.Result set cache can be used between different virtual warehouses

upvoted 2 times

67422cb 2 years, 4 months ago

How is F not a correct answer? It is more correct than B, since a result cache be kicked out of memory before 24 hours if there are newer results coming in.

upvoted 2 times

Balakt 1 year, 8 months ago

Result set cache can be used between different virtual warehouses

upvoted 1 times

carloscloud 2 years, 9 months ago

F, E, C, B, A is valid. I have a doubt with A, but the rest is fine. Are these questions well posed? Or does it refer only to the features that are explicitly in the documentation? B, C, E 100%, since it is in the documentation explicitly, but F AND B are fine implicitly.

upvoted 1 times

GLEM 3 years, 2 months ago

B,C,E, validated

upvoted 2 times

hillcat111 3 years, 3 months ago

Answer is B,C,E and is validated

upvoted 1 times

Needium 3 years, 4 months ago

Chosen answer BCE is correct

upvoted 2 times

serg_khar 3 years, 6 months ago

https://docs.snowflake.com/en/user-guide/querying-persisted-results

upvoted 1 times

Question #10 *Topic 1*

Which organization-related tasks can be performed by the ORGADMIN role? (Choose three.)

A. Changing the name of the organization

B. Creating an account

C. Viewing a list of organization accounts

D. Changing the name of an account

E. Deleting an account

F. Enabling the replication of a database

Suggested Answer: *B, C, F*

Community vote distribution



Jay_98_11 Highly Voted 3 years, 5 months ago

Doc: https://docs.snowflake.com/en/user-guide/organizations#orgadmin-role

upvoted 6 times

uncledana Most Recent 8 months, 1 week ago

Doc: https://docs.snowflake.com/en/user-guide/organizations#orgadmin-role

upvoted 1 times

7b6c707 8 months, 1 week ago

https://docs.snowflake.com/en/user-guide/organizations-manage-accounts

upvoted 1 times

Saai123 9 months ago

Considering without any limitations or restrictions BCD is the right answer

upvoted 1 times

BobFar 1 year, 1 month ago

There are 4 correct answers BCDE Administrators use the organization account to do the following: View organization-level data collected from all accounts in the organization, including the query history from each account. Manage organization users and user groups for all accounts, or a subset of accounts, in an organization. Enable Snowflake Marketplace terms for the entire organization. Manage the lifecycle of accounts in an organization, including creating and deleting accounts. Enable replication for an account. E: Dropping account also has some exception, You cannot drop an account that has active listings shared to specific consumers or listings published on the Snowflake Marketplace.
https://docs.snowflake.com/en/user-guide/organization-accounts A: is wrong need support ticket F: is wrong , can enable replication for the account not for a specific database

upvoted 3 times

qqcvd 1 year, 4 months ago

it is not F because "Account administrators (users with the ACCOUNTADMIN role) can manage replication and failover/failback actions in Snowsight or Classic Console "

upvoted 2 times

nareeshk 1 year, 8 months ago

Incorrect choices. BCDE are possible i.e. create / drop / rename / list accounts. With ORGADMIN one can enable replication for an account, and not for a database, hence F is not valid option

upvoted 2 times

MSIDDIQUI18 1 year, 11 months ago

D is also correct with this commad; ALTER ACCOUNT RENAME TO [SAVE_OLD_URL = { TRUE | FALSE }]

upvoted 2 times

mns0173 2 years, 7 months ago

D and E are account level operations, not organization level A can only be done via support ticket So BCF

upvoted 2 times

jjordan 2 years, 9 months ago

Selected Answer: BCE

BCE is correct.

upvoted 2 times

jjordan 2 years, 9 months ago

BCE is correct.I meant BCF

upvoted 2 times

jjordan 2 years, 9 months ago

I meant BCF

upvoted 1 times

GLEM 3 years, 2 months ago

Selected Answer: BCD

There are 4 valid answer here : A. Changing the name of the organization --> No B. Creating an account --> Yes : https://docs.snowflake.com/en/user-guide/organizations-manage-accounts-create C. Viewing a list of organization accounts --> Yes : https://docs.snowflake.com/en/user-guide/organizations-manage-accounts-view D. Changing the name of an account --> Yes : https://docs.snowflake.com/en/user-guide/organizations-manage-accounts-rename E. Deleting an account --> No (Yes but in preview, so I assume it's not true) : https://docs.snowflake.com/en/user-guide/organizations-manage-accounts-delete F. Enabling the replication of a database --> Yes : https://docs.snowflake.com/en/user-guide/organizations#orgadmin-role

upvoted 4 times

GLEM 3 years, 2 months ago

There are 4 valid answer here : A. Changing the name of the organization --> No B. Creating an account --> Yes : https://docs.snowflake.com/en/user-guide/organizations-manage-accounts-create C. Viewing a list of organization accounts --> Yes : https://docs.snowflake.com/en/user-guide/organizations-manage-accounts-view D. Changing the name of an account --> Yes : https://docs.snowflake.com/en/user-guide/organizations-manage-accounts-rename E. Deleting an account --> No (Yes but in preview, so I assume it's not true) : https://docs.snowflake.com/en/user-guide/organizations-manage-accounts-delete F. Enabling the replication of a database --> Yes : https://docs.snowflake.com/en/user-guide/organizations#orgadmin-roleOrgadmin can enable replication on an account in org, not at an individual database level.. F is incorrectThere is a limitation that you cannot rename an account while you are logged in, which I think eliminates this as a correct answer.That is a limitation but its not that you cannot rename an account with Orgadmin role, so D stands correct in my opinion

upvoted 4 times

nareeshk 1 year, 8 months ago

Orgadmin can enable replication on an account in org, not at an individual database level.. F is incorrect

upvoted 1 times

jjordan 2 years, 9 months ago

There is a limitation that you cannot rename an account while you are logged in, which I think eliminates this as a correct answer.That is a limitation but its not that you cannot rename an account with Orgadmin role, so D stands correct in my opinion

upvoted 1 times

jjordan 2 years, 9 months ago

There is a limitation that you cannot rename an account while you are logged in, which I think eliminates this as a correct answer.

upvoted 1 times

AJ2908 2 years, 6 months ago

That is a limitation but its not that you cannot rename an account with Orgadmin role, so D stands correct in my opinion

upvoted 1 times

hillcat111 3 years, 3 months ago

Answer is B,C,F and is validated

upvoted 3 times

qqcvd 3 years, 3 months ago

i thin it is BCE

upvoted 3 times

Alachramkowa 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/organizations-manage-accounts#changing-account-name agree with BCF D is correct in case when organization has at least 2 accounts

upvoted 1 times

Alachramkowa 3 years, 5 months ago

i would go with BCD

upvoted 1 times

serg_khar 3 years, 6 months ago

it can be check here https://docs.snowflake.com/en/user-guide/organizations#orgadmin-role

upvoted 3 times

henrik0 3 years, 6 months ago

wrong. https://docs.snowflake.com/en/user-guide/organizations#orgadmin-role

upvoted 1 times

Question #11 *Topic 1*

A Data Engineer is designing a near real-time ingestion pipeline for a retail company to ingest event logs into Snowflake to derive insights. A Snowflake Architect is asked to define security best practices to configure access control privileges for the data load for auto-ingest to Snowpipe. What are the MINIMUM object privileges required for the Snowpipe user to execute Snowpipe?

A. OWNERSHIP on the named pipe, USAGE on the named stage, target database, and schema, and INSERT and SELECT on the target table

B. OWNERSHIP on the named pipe, USAGE and READ on the named stage, USAGE on the target database and schema, and INSERT end SELECT on the target table

C. CREATE on the named pipe, USAGE and READ on the named stage, USAGE on the target database and schema, and INSERT end SELECT on the target table

D. USAGE on the named pipe, named stage, target database, and schema, and INSERT and SELECT on the target table

Suggested Answer: *A*

Community vote distribution

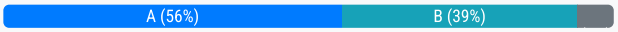

jlad26 Highly Voted 2 years, 7 months ago

Selected Answer: A

Pipe does not have USAGE privilege and internal stage does not have READ privilege so has to be A

upvoted 8 times

jlad26 Highly Voted 2 years, 7 months ago

Pipe does not have USAGE privilege and internal stage does not have READ privilege so has to be ASorry forgot to say pipe does not have CREATE privilege either. https://docs.snowflake.com/en/user-guide/security-access-control-privilegesB is the correct answer https://docs.snowflake.com/user-guide/data-load-snowpipe-auto-s3#step-4-configure-security

upvoted 8 times

jlad26 2 years, 7 months ago

Sorry forgot to say pipe does not have CREATE privilege either. https://docs.snowflake.com/en/user-guide/security-access-control-privileges

upvoted 1 times

BobFar 1 year, 1 month ago

B is the correct answer https://docs.snowflake.com/user-guide/data-load-snowpipe-auto-s3#step-4-configure-security

upvoted 1 times

7b6c707 Most Recent 8 months, 1 week ago

when it says using auto_ingestion, it implies using external stage, therefore, only USAGE on the stage is needed

upvoted 1 times

BobFar 1 year, 1 month ago

Selected Answer: A

A is correct there is no usage for named stage (internal) https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro#owning-pipes

upvoted 1 times

BobFar 1 year, 1 month ago

A is correct there is no usage for named stage (internal) https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro#owning-pipesIt seems I was wrong, B is the corret answer https://docs.snowflake.com/user-guide/data-load-snowpipe-auto-s3#step-4-configure-security

upvoted 1 times

BobFar 1 year, 1 month ago

It seems I was wrong, B is the corret answer https://docs.snowflake.com/user-guide/data-load-snowpipe-auto-s3#step-4-configure-security

upvoted 1 times

785de4c 1 year, 2 months ago

https://docs.snowflake.com/user-guide/data-load-snowpipe-auto-s3#step-4-configure-security

upvoted 3 times

Akhil_8121 2 years ago

option B https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro#creating-pipes Creating pipes Creating and managing pipes requires a role with a minimum of the following privileges: Object Privilege Notes Database USAGE Schema USAGE, CREATE PIPE Stage in the pipe definition USAGE External stages only. Stage in the pipe definition READ Internal stages only. Table in the pipe definition SELECT, INSERT

upvoted 3 times

oracle901 2 years, 4 months ago

To view details about the pipe, you must use a role with the MONITOR or OWNERSHIP privilege on the pipe and the USAGE privilege on both the database and schema that contain the pipe. See Pipe privileges for more details. https://docs.snowflake.com/user-guide/data-load-snowpipe-manage

upvoted 1 times

oracle901 2 years, 4 months ago

To view details about the pipe, you must use a role with the MONITOR or OWNERSHIP privilege on the pipe and the USAGE privilege on both the database and schema that contain the pipe. See Pipe privileges for more details. https://docs.snowflake.com/user-guide/data-load-snowpipe-manageA is the answer

upvoted 1 times

oracle901 2 years, 4 months ago

A is the answer

upvoted 1 times

Vishal27dec 3 years ago

Ans is A:

-- Create a role to contain the Snowpipe privileges use role securityadmin;

create or replace role snowpipe1;

-- Grant the required privileges on the database objects grant usage on database snowpipe_db to role snowpipe1;

grant usage on schema snowpipe_db.public to role snowpipe1;

grant insert, select on snowpipe_db.public.mytable to role snowpipe1;

grant usage on stage snowpipe_db.public.mystage to role snowpipe1;

-- Grant the OWNERSHIP privilege on the pipe object grant ownership on pipe snowpipe_db.public.mypipe to role snowpipe1;

-- Grant the role to a user grant role snowpipe1 to user jsmith;

-- Set the role as the default role for the user alter user jsmith set default_role = snowpipe1;

upvoted 4 times

Vishal27dec 3 years ago

Ans is A: -- Create a role to contain the Snowpipe privileges use role securityadmin; create or replace role snowpipe1; -- Grant the required privileges on the database objects grant usage on database snowpipe_db to role snowpipe1; grant usage on schema snowpipe_db.public to role snowpipe1; grant insert, select on snowpipe_db.public.mytable to role snowpipe1; grant usage on stage snowpipe_db.public.mystage to role snowpipe1; -- Grant the OWNERSHIP privilege on the pipe object grant ownership on pipe snowpipe_db.public.mypipe to role snowpipe1; -- Grant the role to a user grant role snowpipe1 to user jsmith; -- Set the role as the default role for the user alter user jsmith set default_role = snowpipe1;So, it would be an internal stage by the USAGE on stage. This is the answer that makes the most sense, thank you!

upvoted 4 times

carloscloud 2 years, 9 months ago

So, it would be an internal stage by the USAGE on stage. This is the answer that makes the most sense, thank you!

upvoted 1 times

Saai123 3 years, 2 months ago

Answer is A. OWNERSHIP on the named pipe, USAGE on the named stage, target database, and schema, and INSERT and SELECT on the target table

upvoted 3 times

hillcat111 3 years, 3 months ago

Answer is D and is validated

upvoted 1 times

qqcvd 3 years, 3 months ago

it doesn't exist a GRANT USAGE on PIPE https://docs.snowflake.com/en/user-guide/security-access-control-privileges#pipe-privileges As it is about auto-ingest, it's not possible with internal STAGE, so it is not B or C as READ grant is only for internal STAGE Thus, it it response D

upvoted 1 times

Alachramkowa 3 years, 5 months ago

C https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro#creating-pipes

upvoted 1 times

serg_khar 3 years, 6 months ago

https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro#owning-pipes

upvoted 4 times

Question #12 *Topic 1*

The IT Security team has identified that there is an ongoing credential stuffing attack on many of their organization's system.
What is the BEST way to find recent and ongoing login attempts to Snowflake?

A. Call the LOGIN_HISTORY Information Schema table function.

B. Query the LOGIN_HISTORY view in the ACCOUNT_USAGE schema in the SNOWFLAKE database.

C. View the History tab in the Snowflake UI and set up a filter for SQL text that contains the text "LOGIN".

D. View the Users section in the Account tab in the Snowflake UI and review the last login column.

Suggested Answer: *A*

Community vote distribution

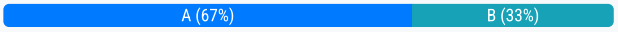

GLEM Highly Voted 3 years, 2 months ago

Key word here : way to find recent and ongoing login attempts . So can't be the B answer, as there is lag (2 or 3 hours) in the account_usage view. Answer is A.

upvoted 8 times

MSIDDIQUI18 Most Recent 1 year, 11 months ago

Selected Answer: A If TIME_RANGE_END is not specified, the function returns the most recent login events. If the time range does not fall within the last 7 days, an error is returned. select * from table(information_schema.login_history_by_user()) order by event_timestamp; select * from table(information_schema.login_history(TIME_RANGE_START => dateadd('hours',-1,current_timestamp()),current_timestamp())) order by event_timestamp;

upvoted 1 times

jjordan 2 years, 9 months ago

Selected Answer: B

https://docs.snowflake.com/en/user-guide/data-load-snowpipe-auto-s3#step-3-configure-security

upvoted 1 times

jjordan 2 years, 9 months ago

https://docs.snowflake.com/en/user-guide/data-load-snowpipe-auto-s3#step-3-configure-securityShould be A

upvoted 1 times

jjordan 2 years, 9 months ago

Should be A

upvoted 1 times

hillcat111 3 years, 3 months ago

Answer is A and is validated

upvoted 1 times

victorleonis 3 years, 4 months ago

LOGIN_HISTORY is not in the information schema it's part of ACCOUNT_USAGE https://docs.snowflake.com/en/sql-reference/account-usage/login_historyKey word here : way to find recent and ongoing login attempts . So can't be the B answer, as there is lag (2 or 3 hours) in the account_usage view. Answer is A.the most up to date data is here : https://docs.snowflake.com/en/sql-reference/functions/login_history so it is response A

upvoted 3 times

victorleonis 3 years, 4 months ago

Selected Answer: B

LOGIN_HISTORY is not in the information schema it's part of ACCOUNT_USAGE https://docs.snowflake.com/en/sql-reference/account-usage/login_history

upvoted 3 times

qqcvd 3 years, 3 months ago

the most up to date data is here : https://docs.snowflake.com/en/sql-reference/functions/login_history so it is response A

upvoted 2 times

Question #13 *Topic 1*

An Architect has a VPN_ACCESS_LOGS table in the SECURITY_LOGS schema containing timestamps of the connection and disconnection, username of the user, and summary statistics.
What should the Architect do to enable the Snowflake search optimization service on this table?

A. Assume role with OWNERSHIP on future tables and ADD SEARCH OPTIMIZATION on the SECURITY_LOGS schema.

B. Assume role with ALL PRIVILEGES including ADD SEARCH OPTIMIZATION in the SECURITY LOGS schema.

C. Assume role with OWNERSHIP on VPN_ACCESS_LOGS and ADD SEARCH OPTIMIZATION in the SECURITY_LOGS schema.

D. Assume role with ALL PRIVILEGES on VPN_ACCESS_LOGS and ADD SEARCH OPTIMIZATION in the SECURITY_LOGS schema.

Suggested Answer: *C*

Community vote distribution

C (100%)

HighBMI Highly Voted 2 years, 4 months ago

https://docs.snowflake.com/en/user-guide/search-optimization-service 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.

upvoted 7 times

hillcat111 Highly Voted 2 years, 3 months ago

Answer is C and is validated

upvoted 5 times

DhirajG Most Recent 1 year, 6 months ago

Ownership is required on table on which SOS needs to be added.

upvoted 1 times

Question #14 *Topic 1*

The table contains five columns and it has millions of records. The cardinality distribution of the columns is shown below:

Column C4 and C5 are mostly used by SELECT queries in the GROUP BY and ORDER BY clauses. Whereas columns C1, C2 and C3 are heavily used in filter and join conditions of SELECT queries.
The Architect must design a clustering key for this table to improve the query performance.
Based on Snowflake recommendations, how should the clustering key columns be ordered while defining the multi-column clustering key?

Column	Number of Distinct Values
C1	10,790
C2	108
C3	302,605
C4	1.117,736
C5	2.205,400

A. C5, C4, C2

B. C3, C4, C5

C. C1, C3, C2

D. C2, C1, C3

Suggested Answer: *D*

Community vote distribution

D (100%)

Jay_98_11 Highly Voted 2 years, 11 months ago

correct

upvoted 7 times

caio_tertuliano Most Recent 1 year, 7 months ago

Question D. Saw this exactly on the actual exam.

upvoted 1 times

Balakt 1 year, 8 months ago

C4 and C5 are ruled out. C2, C1 and C3 are correct (Based on the increasing cardinality)

upvoted 1 times

jlad26 2 years, 1 month ago

If you are defining a multi-column clustering key for a table, the order in which the columns are specified in the CLUSTER BY clause is important. As a general rule, Snowflake recommends ordering the columns from lowest cardinality to highest cardinality. Putting a higher cardinality column before a lower cardinality column will generally reduce the effectiveness of clustering on the latter column.

upvoted 1 times

hillcat111 2 years, 9 months ago

Answer is D and is validated

upvoted 4 times

Question #15 *Topic 1*

Which security, governance, and data protection features require, at a MINIMUM, the Business Critical edition of Snowflake? (Choose two.)

A. Extended Time Travel (up to 90 days)

B. Customer-managed encryption keys through Tri-Secret Secure

C. Periodic rekeying of encrypted data

D. AWS, Azure, or Google Cloud private connectivity to Snowflake

E. Federated authentication and SSO

Suggested Answer: *B, D*

Community vote distribution

BD (100%)

HighBMI Highly Voted 2 years, 4 months ago

B: https://docs.snowflake.com/en/user-guide/security-encryption-manage#tri-secret-secure D: https://docs.snowflake.com/en/user-guide/admin-security-privatelink

upvoted 5 times

Balakt Most Recent 1 year, 8 months ago

A, C - Enterprise and above E - For all Editions Hence B and E are correct

upvoted 1 times

VenkataSri 1 year, 10 months ago

Answer is B,D and is validated

upvoted 1 times

hillcat111 2 years, 3 months ago

Answer is B,D and is validated

upvoted 2 times

Question #16 *Topic 1*

A company wants to deploy its Snowflake accounts inside its corporate network with no visibility on the internet. The company is using a VPN infrastructure and Virtual Desktop Infrastructure (VDI) for its Snowflake users. The company also wants to re-use the login credentials set up for the VDI to eliminate redundancy when managing logins.
What Snowflake functionality should be used to meet these requirements? (Choose two.)

A. Set up replication to allow users to connect from outside the company VPN.

B. Provision a unique company Tri-Secret Secure key.

C. Use private connectivity from a cloud provider.

D. Set up SSO for federated authentication.

E. Use a proxy Snowflake account outside the VPN, enabling client redirect for user logins.

Suggested Answer: *C, D*

Community vote distribution

CD (100%)

MSIDDIQUI18 Most Recent 1 year, 11 months ago
Correct Answer: CD, To the Snowflake Service When the routing is through a private IP address from your VPC or VNET to the Snowflake VPC or VNet, that is private connectivity to the Snowflake Service. These connections use AWS Private Link, Azure Private Link, or Google Cloud Private Service Connect. The service depends on the cloud platform that hosts your Snowflake account. To internal stages Similarly, you can use private connectivity to connect to Snowflake internal stages for accounts on AWS and Azure. This is private connectivity to Snowflake internal stages. To external services Available to all accounts.
upvoted 2 times

DBSS 2 years, 5 months ago
BC, Tri-Secret and VPN are keyI meant BD
upvoted 1 times

DBSS 2 years, 5 months ago
BC, Tri-Secret and VPN are key
upvoted 1 times

DBSS 2 years, 5 months ago
I meant BD
upvoted 1 times

Shreyas1998NewFact 2 years, 5 months ago
CD is correct
upvoted 1 times

hitman000000 2 years, 8 months ago
CD is correct
upvoted 2 times

VenkataSri 2 years, 10 months ago
Answer is C,D and is validated
upvoted 2 times

MariusdeWith 3 years, 1 month ago
CD is correct. Tri-Secret Secure key is for data-encryption and not for securing logins.
upvoted 2 times

GLEM 3 years, 2 months ago
CD are correct for the requirements
upvoted 1 times

hillcat111 3 years, 3 months ago
Answer is C,D and is validated
upvoted 2 times

RAVIRAJA 3 years, 3 months ago

CD looks right

upvoted 2 times

ManojRaghuwanshi 3 years, 3 months ago

CD looks correct to me

upvoted 1 times

SrideviJ 3 years, 4 months ago

I think CD are correct options

upvoted 4 times

Question #17 *Topic 1*

How do Snowflake databases that are created from shares differ from standard databases that are not created from shares? (Choose three.)

A. Shared databases are read-only.

B. Shared databases must be refreshed in order for new data to be visible.

C. Shared databases cannot be cloned.

D. Shared databases are not supported by Time Travel.

E. Shared databases will have the PUBLIC or INFORMATION_SCHEMA schemas without explicitly granting these schemas to the share.

F. Shared databases can also be created as transient databases.

Suggested Answer: *A, C, D*

Community vote distribution

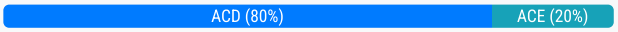

Jay_98_11 Highly Voted 2 years, 11 months ago

Should be ACD

upvoted 5 times

hawkhunt Most Recent 1 year, 8 months ago

https://docs.snowflake.com/en/sql-reference/sql/create-database#general-usage-notes ACD * Databases created from shares differ from standard databases in the following ways: * They do not have the PUBLIC or INFORMATION_SCHEMA schemas unless these schemas were explicitly granted to the share. * They cannot be cloned. * Properties, such as TRANSIENT and DATA_RETENTION_TIME_IN_DAYS, do not apply. In question, the statement says: E. Shared databases will have the PUBLIC or INFORMATION_SCHEMA schemas without explicitly granting these schemas to the share. However, the fact mentioned is: They do not have the PUBLIC or INFORMATION_SCHEMA schemas unless these schemas were explicitly granted to the share. E is not correct.

upvoted 1 times

jjordan 2 years, 4 months ago

https://docs.snowflake.com/en/sql-reference/sql/create-database#general-usage-notes - Databases created from shares differ from standard databases in the following ways: - They do not have the PUBLIC or INFORMATION_SCHEMA schemas unless these schemas were explicitly granted to the share. - They cannot be cloned. Properties, such as TRANSIENT and DATA_RETENTION_TIME_IN_DAYS, do not apply.So why E?

upvoted 2 times

jjordan 2 years, 4 months ago

Selected Answer: ACE

https://docs.snowflake.com/en/sql-reference/sql/create-database#general-usage-notes

- Databases created from shares differ from standard databases in the following ways: - They do not have the PUBLIC or INFORMATION_SCHEMA schemas unless these schemas were explicitly granted to the share. - They cannot be cloned.

Properties, such as TRANSIENT and DATA_RETENTION_TIME_IN_DAYS, do not apply.

upvoted 2 times

umidjon03 2 years, 4 months ago

So why E?

upvoted 1 times

VenkataSri 2 years, 4 months ago

Answer ACD is correct

upvoted 2 times

hillcat111 2 years, 9 months ago

Answer is A,D,E and is validated Explanation 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.

Databases created from shares differ from standard databases in the following ways:

They do not have the PUBLIC or INFORMATION_SCHEMA schemas unless these schemas were explicitly granted to the share.

They cannot be cloned.

Properties, such as TRANSIENT and DATA_RETENTION_TIME_IN_DAYS, do not apply.

upvoted 2 times

hillcat111 2 years, 9 months ago

Answer is A,D,E and is validated Explanation 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. Databases created from shares differ from standard databases in the following ways: They do not have the PUBLIC or INFORMATION_SCHEMA schemas unless these schemas were explicitly granted to the share. They cannot be cloned. Properties, such as TRANSIENT and DATA_RETENTION_TIME_IN_DAYS, do not apply.These explanation indicate answer ACD though.So answer should be ACD right...

upvoted 2 times

GLEM 2 years, 8 months ago

These explanation indicate answer ACD though.

upvoted 3 times

prshntdxt7 2 years, 1 month ago

So answer should be ACD right...

upvoted 1 times

serg_khar 2 years, 11 months ago

User qqcvd provided explanation

upvoted 2 times

qqcvd 3 years ago

The following actions are not supported: Time Travel for a shared database or any schemas/tables in the database. -> https://docs.snowflake.com/en/user-guide/data-share-consumers Databases created from shares differ from standard databases in the following ways: They do not have the PUBLIC or INFORMATION_SCHEMA schemas unless these schemas were explicitly granted to the share. -> https://docs.snowflake.com/en/sql-reference/sql/create-database

upvoted 1 times

Question #18 *Topic 1*

What integration object should be used to place restrictions on where data may be exported?

A. Stage integration

B. Security integration

C. Storage integration

D. API integration

Suggested Answer: *C*

Community vote distribution

C (100%)

GLEM Most Recent 2 years, 2 months ago

Storage integration : https://docs.snowflake.com/en/sql-reference/sql/create-storage-integration

upvoted 3 times

hillcat111 2 years, 3 months ago

Answer is C and is validated

upvoted 1 times

hillcat111 2 years, 4 months ago

Sorry answer C

upvoted 1 times

hillcat111 2 years, 4 months ago

I think this is B

upvoted 1 times

Question #19 *Topic 1*

The following DDL command was used to create a task based on a stream:

Assuming MY_WH is set to auto_suspend – 60 and used exclusively for this task, which statement is true?

```
CREATE TASK ts_insert_new_customers
  WAREHOUSE = MY_WH
  Schedule = '5 minute'
WHEN
  System$STREAM_HAS_DATA('MYSTREAM')
AS
  INSERT INTO new customers(id, name) SELECT id, name
FROM mystream WHERE METADATA$ACTION = 'INSERT';
```

A. The warehouse MY_WH will be made active every five minutes to check the stream.

B. The warehouse MY_WH will only be active when there are results in the stream.

C. The warehouse MY_WH will never suspend.

D. The warehouse MY_WH will automatically resize to accommodate the size of the stream.

Suggested Answer: *B*

Community vote distribution

B (100%)

Jay_98_11 Highly Voted 2 years, 5 months ago

Should be B

upvoted 6 times

jjordan Most Recent 1 year, 10 months ago

B is correct.

upvoted 1 times

GLEM 2 years, 2 months ago

Stream_has_data is a metadata operation, so you don't need a wh to call it. B is correct, the WH will start as soon as there are data in the STREAM

upvoted 4 times

hillcat111 2 years, 3 months ago

Answer is B and is validated

upvoted 2 times

victorleonis 2 years, 4 months ago

the warehouse will be triggered if condition STREAM_HAS_DATA meet

upvoted 3 times

taod 2 years, 5 months ago

Not true, B would be correct. Warehouse is only started if condition is met: https://docs.snowflake.com/en/sql-reference/sql/create-task
https://docs.snowflake.com/en/sql-reference/functions/system_stream_has_data

upvoted 3 times

Question #20 *Topic 1*

What is a characteristic of loading data into Snowflake using the Snowflake Connector for Kafka?

A. The Connector only works in Snowflake regions that use AWS infrastructure.

B. The Connector works with all file formats, including text, JSON, Avro, Ore, Parquet, and XML.

C. The Connector creates and manages its own stage, file format, and pipe objects.

D. Loads using the Connector will have lower latency than Snowpipe and will ingest data in real time.

Suggested Answer: *C*

Community vote distribution

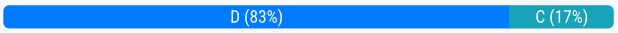

7b6c707 Most Recent 8 months, 1 week ago

snowflake kafka connector (either using snowpipe or snowpipe streaming) won't need file_format, as it only deals with json or avro file type. these semi structure files get load into variant column

upvoted 2 times

DhirajG 1 year, 6 months ago

Snowpipe Streaming using Kafka load the data in real time which has latency lesser than snowpipe. Correct Answer is D

upvoted 1 times

ywan1600 1 year, 7 months ago

C is the answer

upvoted 1 times

nareeshk 1 year, 8 months ago

How is this Answer C? Kafka doesnt use a file format. It only uses internal stage, pipe & table

upvoted 2 times

hillcat111 2 years, 3 months ago

Answer is C and is validated

upvoted 3 times

hillcat111 2 years, 4 months ago

I think C is the right answer

upvoted 2 times

Question #21 *Topic 1*

A healthcare company wants to share data with a medical institute. The institute is running a Standard edition of Snowflake; the healthcare company is running a Business Critical edition.
How can this data be shared?

A. The healthcare company will need to change the institute's Snowflake edition in the accounts panel.

B. By default, sharing is supported from a Business Critical Snowflake edition to a Standard edition.

C. Contact Snowflake and they will execute the share request for the healthcare company.

D. Set the share_restriction parameter on the shared object to false.

Suggested Answer: *D*

Community vote distribution



serg_khar Highly Voted 3 years, 6 months ago
https://docs.snowflake.com/en/user-guide/override_share_restrictions For me it looks as D: -- grant the privilege to the SYSADMIN role use role accountadmin; grant override share restrictions on account to role sysadmin; -- SYSADMIN can now add a consumer account to a share with the SHARE_RESTRICTIONS parameter set to false use role sysadmin; alter share add accounts = SHARE_RESTRICTIONS=false;
upvoted 5 times

Balakt Most Recent 1 year, 8 months ago
B->Not by default C-> If it is not working then only one can contact D ->is correct.
upvoted 1 times

laksnarn 2 years, 3 months ago
B. By default, sharing is supported from a Business Critical Snowflake edition to a Standard edition. Snowflake allows data sharing between different editions of accounts, including from higher editions (like Business Critical) to lower editions (like Standard). This means that the healthcare company can share data with the medical institute without needing to change the institute's Snowflake edition or contact Snowflake for assistance.
upvoted 3 times

MariusdeWith 2 years, 11 months ago
It is D. https://docs.snowflake.com/en/user-guide/override_share_restrictions -> By default, Snowflake does not allow sharing data from a Business Critical to a non-Business Critical account.
upvoted 3 times

PedroSDG 2 years, 11 months ago
Not D. SHARE_RESTRICTIONS is an account parameter, not an object
upvoted 1 times

PedroSDG 2 years, 11 months ago
It is B. Sharing is SUPPORTED from business critical to any, but not ENABLED.
upvoted 2 times

hillcat111 3 years, 3 months ago
Answer is D and is validated
upvoted 1 times

Question #22 *Topic 1*

An Architect is designing a pipeline to stream event data into Snowflake using the Snowflake Kafka connector. The Architect's highest priority is to configure the connector to stream data in the MOST cost-effective manner.
Which of the following is recommended for optimizing the cost associated with the Snowflake Kafka connector?

A. Utilize a higher Buffer.flush.time in the connector configuration.

B. Utilize a higher Buffer.size.bytes in the connector configuration.

C. Utilize a lower Buffer.size.bytes in the connector configuration.

D. Utilize a lower Buffer.count.records in the connector configuration.

Suggested Answer: *A*

Community vote distribution

A (100%)

victorleonis Highly Voted 2 years, 4 months ago

The minimum value supported for the buffer.flush.time property is 1 (in seconds). For higher average data flow rates, we suggest that you decrease the default value for improved latency. If cost is a greater concern than latency, you could increase the buffer flush time. Be careful to flush the Kafka memory buffer before it becomes full to avoid out of memory exceptions. https://docs.snowflake.com/en/user-guide/data-load-snowpipe-streaming-kafka

upvoted 5 times

jjordan Most Recent 1 year, 10 months ago

https://docs.snowflake.com/en/user-guide/data-load-snowpipe-streaming-kafka

upvoted 2 times

hillcat111 2 years, 3 months ago

Answer is A and is validated

upvoted 2 times

Alachramkowa 2 years, 5 months ago

With lower buffer.count.records the buffer limits will be reached earlier, the file will be flushed and sent for ingestion through Snowpipe more frequently, thus increasing the cost

upvoted 3 times

Question #23 *Topic 1*

An Architect has chosen to separate their Snowflake Production and QA environments using two separate Snowflake accounts.
The QA account is intended to run and test changes on data and database objects before pushing those changes to the Production account. It is a requirement that all database objects and data in the QA account need to be an exact copy of the database objects, including privileges and data in the Production account on at least a nightly basis.
Which is the LEAST complex approach to use to populate the QA account with the Production account's data and database objects on a nightly basis?

A. 1. Create a share in the Production account for each database 2. Share access to the QA account as a Consumer 3. The QA account creates a database directly from each share 4. Create clones of those databases on a nightly basis 5. Run tests directly on those cloned databases

B. 1. Create a stage in the Production account 2. Create a stage in the QA account that points to the same external object-storage location 3. Create a task that runs nightly to unload each table in the Production account into the stage 4. Use Snowpipe to populate the QA account

C. 1. Enable replication for each database in the Production account 2. Create replica databases in the QA account 3. Create clones of the replica databases on a nightly basis 4. Run tests directly on those cloned databases

D. 1. In the Production account, create an external function that connects into the QA account and returns all the data for one specific table 2. Run the external function as part of a stored procedure that loops through each table in the Production account and populates each table in the QA account

Suggested Answer: *C*

Community vote distribution

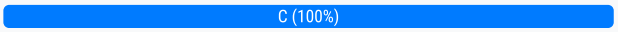

Yogendra_examtopic Most Recent 1 year, 8 months ago
It's simple Question needs same privileges as of production Share dosnt provide it but replication provide it So easiest way is replicate not share
upvoted 1 times

nareeshk 1 year, 8 months ago
Lots of issues and gaps in question & options provided. There is no mention of region for prod and dev account, which dictates the choice between Share & replication.
upvoted 1 times

NachoPrendes 2 years, 2 months ago
But it is not "shared", is "replicated". And we can create a clone from a replicated database. It's C
upvoted 1 times

rkarthik0789 2 years, 4 months ago
The least complex approach among the options provided is likely option B. This approach involves setting up a shared stage between the Production and QA accounts and utilizing Snowpipe for data ingestion. Snowpipe can automatically load data from the stage into the QA environment, simplifying the process and reducing the need for manual intervention. Additionally, unloading data from the Production account into a stage is a straightforward task, and Snowpipe's automated data loading capabilities streamline the process further.
upvoted 1 times

anikl 2 years, 11 months ago
Share database cannot be cloned. therefore A is wrong, C is correct
upvoted 2 times

hillcat111 3 years, 3 months ago
Answer is C and is validated
upvoted 1 times

victorleonis 3 years, 4 months ago
Ans C is incorrect. We can clone a shared object.
upvoted 2 times

Dhamod 3 years, 4 months ago
https://community.snowflake.com/s/article/Snowflake-DTAP-Environment-Setup Ans: C
upvoted 2 times

callipso21 3 years, 5 months ago

https://community.snowflake.com/s/article/How-to-Copy-a-Database-from-One-account-to-Another

upvoted 4 times

Question #24 *Topic 1*

A user can change object parameters using which of the following roles?

A. ACCOUNTADMIN, SECURITYADMIN

B. SYSADMIN, SECURITYADMIN

C. ACCOUNTADMIN, USER with PRIVILEGE

D. SECURITYADMIN, USER with PRIVILEGE

Suggested Answer: *C*

Community vote distribution

C (100%)

DBSS Most Recent 1 year, 11 months ago

C for sure

upvoted 1 times

hillcat111 2 years, 9 months ago

Answer is C and is validated

upvoted 2 times

Jay_98_11 2 years, 11 months ago

should be C

upvoted 2 times

callipso21 2 years, 11 months ago

Object Parameters¶ Object parameters can be set at the following levels: Account Account administrators can use the ALTER ACCOUNT command to set object parameters for the account. The values set for the account default to the objects created in the account. Object Users with the appropriate privileges can use the corresponding CREATE or ALTER commands to override object parameters for an individual object.

upvoted 3 times

serg_khar 2 years, 11 months ago

https://docs.snowflake.com/en/sql-reference/parameters#object-parameters

upvoted 4 times

Question #25 *Topic 1*

A media company needs a data pipeline that will ingest customer review data into a Snowflake table, and apply some transformations. The company also needs to use Amazon Comprehend to do sentiment analysis and make the de-identified final data set available publicly for advertising companies who use different cloud providers in different regions.
The data pipeline needs to run continuously ang efficiently as new records arrive in the object storage leveraging event notifications. Also, the operational complexity, maintenance of the infrastructure, including platform upgrades and security, and the development effort should be minimal.
Which design will meet these requirements?

A. Ingest the data using COPY INTO and use streams and tasks to orchestrate transformations. Export the data into Amazon S3 to do model inference with Amazon Comprehend and ingest the data back into a Snowflake table. Then create a listing in the Snowflake Marketplace to make the data available to other companies.

B. Ingest the data using Snowpipe and use streams and tasks to orchestrate transformations. Create an external function to do model inference with Amazon Comprehend and write the final records to a Snowflake table. Then create a listing in the Snowflake Marketplace to make the data available to other companies.

C. Ingest the data into Snowflake using Amazon EMR and PySpark using the Snowflake Spark connector. Apply transformations using another Spark job. Develop a python program to do model inference by leveraging the Amazon Comprehend text analysis API. Then write the results to a Snowflake table and create a listing in the Snowflake Marketplace to make the data available to other companies.

D. Ingest the data using Snowpipe and use streams and tasks to orchestrate transformations. Export the data into Amazon S3 to do model inference with Amazon Comprehend and ingest the data back into a Snowflake table. Then create a listing in the Snowflake Marketplace to make the data available to other companies.

Suggested Answer: *B*

Community vote distribution

B (100%)

hillcat111 Most Recent 2 years, 3 months ago
Answer is B and is validated
upvoted 2 times

HighBMI 2 years, 4 months ago
https://aws.amazon.com/blogs/apn/using-amazon-comprehend-medical-with-the-snowflake-data-cloud/
upvoted 1 times

Question #26 *Topic 1*

A Snowflake Architect is designing an application and tenancy strategy for an organization where strong legal isolation rules as well as multi-tenancy are requirements.
Which approach will meet these requirements if Role-Based Access Policies (RBAC) is a viable option for isolating tenants?

A. Create accounts for each tenant in the Snowflake organization.

B. Create an object for each tenant strategy if row level security is viable for isolating tenants.

C. Create an object for each tenant strategy if row level security is not viable for isolating tenants.

D. Create a multi-tenant table strategy if row level security is not viable for isolating tenants.

Suggested Answer: *B*

Community vote distribution

C (100%)

ArunH Most Recent 10 months, 1 week ago
If RBAC is a viable option for separation then isolation can be done using Object Per Tenant (OPT).
upvoted 1 times

MSIDDIQUI18 1 year, 11 months ago
Answer "A", according to these two points from this link. • Isolates tenants, thereby reducing the risk of mismanaging security. • Allows for strict security measures (encryption keys, IP allow lists, better-than-RBAC controls) by isolating tenants by account.The question states RBAC is viable, Roles are under the account
upvoted 1 times

MSIDDIQUI18 1 year, 11 months ago
Answer "A", according to these two points from this link. • Isolates tenants, thereby reducing the risk of mismanaging security. • Allows for strict security measures (encryption keys, IP allow lists, better-than-RBAC controls) by isolating tenants by account.
upvoted 1 times

IndunaZulu 1 year, 7 months ago
The question states RBAC is viable, Roles are under the account
upvoted 1 times

SilversRayleigh 2 years, 5 months ago
Why not A? This approach utilizes separate Snowflake accounts for each tenant, ensuring the highest level of isolation as required by strong legal rules, while still allowing multi-tenancy within the broader organizational context. Each tenant's data and processes remain fully isolated in their dedicated account, aligned with stringent compliance needs.
upvoted 2 times

mrango 3 years, 1 month ago
Answer c
upvoted 2 times

hillcat111 3 years, 3 months ago
Answer is C and is validatedbut why not B, as the question says that isolating the tenants is a viable option.You can use a single object for multiple tenants if row-level security is viable.
upvoted 2 times

hillcat111 3 years, 3 months ago
Answer is C and is validated
upvoted 2 times

Soumak 1 year, 8 months ago
but why not B, as the question says that isolating the tenants is a viable option.You can use a single object for multiple tenants if row-level security is viable.
upvoted 1 times

Soumak 1 year, 8 months ago
but why not B, as the question says that isolating the tenants is a viable option.
upvoted 1 times

ywan1600 1 year, 7 months ago

You can use a single object for multiple tenants if row-level security is viable.

upvoted 1 times

Alachramkowa 3 years, 5 months ago

C https://resources.snowflake.com/white-paper/design-patterns-for-building-multi-tenant-applications-on-snowflake

upvoted 1 times

Question #27 *Topic 1*

Which statements describe characteristics of the use of materialized views in Snowflake? (Choose two.)

A. They can include ORDER BY clauses.

B. They cannot include nested subqueries.

C. They can include context functions, such as CURRENT_TIME().

D. They can support MIN and MAX aggregates.

E. They can support inner joins, but not outer joins.

Suggested Answer: *B, D*

Community vote distribution

BD (100%)

hillcat111 Most Recent 2 years, 3 months ago

Answer is B,D and is validated

upvoted 1 times

ManojRaghuwanshi 2 years, 4 months ago

Time functions are non deterministic

upvoted 1 times

victorleonis 2 years, 4 months ago

https://docs.snowflake.com/en/user-guide/views-materialized

upvoted 2 times

serg_khar 2 years, 6 months ago

https://docs.snowflake.com/en/user-guide/views-materialized#limitations-on-creating-materialized-views * The aggregate functions that are supported in materialized views are: MIN/MAX * Functions used in a materialized view must be deterministic. For example, using CURRENT_TIME or CURRENT_TIMESTAMP is not permitted.

upvoted 1 times

Question #28 *Topic 1*

The Data Engineering team at a large manufacturing company needs to engineer data coming from many sources to support a wide variety of use cases and data consumer requirements which include:
1. Finance and Vendor Management team members who require reporting and visualization
2. Data Science team members who require access to raw data for ML model development
3. Sales team members who require engineered and protected data for data monetization
What Snowflake data modeling approaches will meet these requirements? (Choose two.)

A. Consolidate data in the company's data lake and use EXTERNAL TABLES.

B. Create a raw database for landing and persisting raw data entering the data pipelines.

C. Create a set of profile-specific databases that aligns data with usage patterns.

D. Create a single star schema in a single database to support all consumers' requirements.

E. Create a Data Vault as the sole data pipeline endpoint and have all consumers directly access the Vault.

Suggested Answer: *D, E*

Community vote distribution

BC (100%)

hillcat111 Highly Voted 2 years, 9 months ago
Answer is C,E and is validatedHaving all users directly access it may not be ideal. Data Vault structures can be intricate, making direct access challenging for end-users who require simplified and curated datasets for their analyses.
upvoted 7 times

hillcat111 Highly Voted 2 years, 9 months ago
Answer is C,E and is validated
upvoted 7 times

ywan1600 1 year, 7 months ago
Having all users directly access it may not be ideal. Data Vault structures can be intricate, making direct access challenging for end-users who require simplified and curated datasets for their analyses.
upvoted 1 times

Saai123 Most Recent 9 months ago
As raw data is needed by Data Science team & each team has different usecase of data requirements it's not possible to fit into a single model
upvoted 1 times

67422cb 1 year, 10 months ago
A Data Vault is very expansive. Selecting solely from the Data Vault can include a whole bunch of differently structured tables.
upvoted 1 times

ayushi7rawat 2 years, 3 months ago
I think options are BC
upvoted 2 times

SrideviJ 2 years, 10 months ago
I think options are BC
upvoted 3 times

Question #29 *Topic 1*

An Architect on a new project has been asked to design an architecture that meets Snowflake security, compliance, and governance requirements as follows:
1. Use Tri-Secret Secure in Snowflake
2. Share some information stored in a view with another Snowflake customer
3. Hide portions of sensitive information from some columns
4. Use zero-copy cloning to refresh the non-production environment from the production environment
To meet these requirements, which design elements must be implemented? (Choose three.)

A. Define row access policies.

B. Use the Business Critical edition of Snowflake.

C. Create a secure view.

D. Use the Enterprise edition of Snowflake.

E. Use Dynamic Data Masking.

F. Create a materialized view.

Suggested Answer: *B, C, E*

Community vote distribution

BCE (100%)

Saai123 Most Recent 9 months ago
BCE is correct
upvoted 1 times

meng 2 years, 2 months ago
Use Tri-Secret Secure in Snowflake -> Use the Business Critical edition Share some information stored in a view with another Snowflake customer -> Secure View Hide portions of sensitive information from some columns-> Use Dynamic Data Masking. Use zero-copy cloning to refresh the non-production environment from the production environment -> Secure View
upvoted 1 times

jjordan 2 years, 9 months ago
BCE is correct
upvoted 1 times

hillcat111 3 years, 3 months ago
Answer is B,C,E and is validated
upvoted 1 times

callipso21 3 years, 5 months ago
BCE is the correct answer
upvoted 2 times

serg_khar 3 years, 6 months ago
F is incorrect, because you can share only secured objects: https://docs.snowflake.com/en/user-guide/data-sharing-intro
upvoted 3 times

Question #30 *Topic 1*

Which of the following are characteristics of how row access policies can be applied to external tables? (Choose three.)

A. An external table can be created with a row access policy, and the policy can be applied to the VALUE column.

B. A row access policy can be applied to the VALUE column of an existing external table.

C. A row access policy cannot be directly added to a virtual column of an external table.

D. External tables are supported as mapping tables in a row access policy.

E. While cloning a database, both the row access policy and the external table will be cloned.

F. A row access policy cannot be applied to a view created on top of an external table.

Suggested Answer: *A, B, C*

Community vote distribution

ABC (100%)

hillcat111 Most Recent 2 years, 3 months ago

Answer is A,B,C and is validated

upvoted 1 times

callipso21 2 years, 5 months ago

https://docs.snowflake.com/en/user-guide/security-row-intro#external-tables External Tables You can create an external table with a row access policy by executing a CREATE EXTERNAL TABLE statement and apply the policy to the VALUE column. You can apply the row access policy to VALUE column of an existing external table by executing an ALTER TABLE statement on the external table. A row access policy cannot be added to a virtual column directly. Instead, create a view on the external table and apply the row access policy to the columns on the view.

upvoted 1 times

serg_khar 2 years, 6 months ago

https://docs.snowflake.com/en/user-guide/security-row-intro#external-tables

upvoted 1 times

Question #31 *Topic 1*

An Architect needs to allow a user to create a database from an inbound share.
To meet this requirement, the user's role must have which privileges? (Choose two.)

A. IMPORT SHARE;

B. IMPORT PRIVILEGES;

C. CREATE DATABASE;

D. CREATE SHARE;

E. IMPORT DATABASE;

Suggested Answer: *A, C*

Community vote distribution

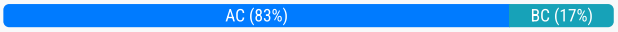

jjordan Most Recent 2 years, 3 months ago
AC is correct.
upvoted 1 times

hillcat111 2 years, 9 months ago
Answer is A,C and is validated
upvoted 1 times

victorleonis 2 years, 10 months ago
IMPORT SHARE Privilege If the IMPORT SHARE privilege is granted to a role, any user with the role can perform the following tasks: View all INBOUND shares (shared by provider accounts) and create databases for the shares. View all OUTBOUND shares owned by the role.
upvoted 2 times

serg_khar 2 years, 11 months ago
https://docs.snowflake.com/en/user-guide/data-share-consumers#step-1-create-a-database-from-the-share Create a database from the share using CREATE DATABASE … FROM SHARE. Executing this command requires a role with the global CREATE DATABASE and IMPORT SHARE privileges.
upvoted 2 times

callipso21 2 years, 12 months ago
https://docs.snowflake.com/en/user-guide/data-share-consumers Allow users to access objects in a share by granting the IMPORTED PRIVILEGES privilege on a shared database to one or more roles in your account.
upvoted 1 times

Question #32 *Topic 1*

Files arrive in an external stage every 10 seconds from a proprietary system. The files range in size from 500 K to 3 MB. The data must be accessible by dashboards as soon as it arrives.
How can a Snowflake Architect meet this requirement with the LEAST amount of coding? (Choose two.)

A. Use Snowpipe with auto-ingest.

B. Use a COPY command with a task.

C. Use a materialized view on an external table.

D. Use the COPY INTO command.

E. Use a combination of a task and a stream.

Suggested Answer: *A, E*

Community vote distribution



Saai123 Most Recent 9 months ago
A & C are correct The input is an external stage A -> is valid option to make a newly arrived files data accessible instantly B-> Copy command with a task scheudle has a lag so not a valid option c-> Valid as external tables + materiablized view combination makes it instantly accesible & performance efficient d-> Just copy into statment manually is not feeasible e-> Stream can't be applied on an external stage it needs an external table so not correct So A & C are correct options
upvoted 2 times

ArunH 10 months, 1 week ago
A & B are possible options independent of each other. We dont need combination of tasks and streams, as there is no use case in this requirement for them.
upvoted 1 times

1940f9d 1 year, 1 month ago
Copy command with task will do . Since it's a file and Option A is by default one.
upvoted 1 times

MSIDDIQUI18 1 year, 11 months ago
According to following statement "C" is correct, but I am not 100% sure, if this logic applied here: "In many cases, the Materialized view on an external table can provide performance that is faster than equivalent queries over the underline External table. "
upvoted 1 times

KindleAce 1 year, 11 months ago
AC Task is using cronjob syntax, which means minimum interval is 60 seconds
upvoted 1 times

h8er 2 years, 5 months ago
A&E is my guess
upvoted 3 times

chiiiik 2 years, 7 months ago
Selected Answer: AE
A is definitely valid Correct me if I'm wrong. E is also valid because if we set STREAM_HAS_DATA it will trigger the task to run to insert data although schedule has been set.
upvoted 2 times

chiiiik 2 years, 7 months ago
A is definitely valid Correct me if I'm wrong. E is also valid because if we set STREAM_HAS_DATA it will trigger the task to run to insert data although schedule has been set.I think this is also true but the question says: with the LEAST amount of codingSorry, I think STREAM_HAS_DATA can not change the scheduling of a task: This function is intended to be used in the WHEN expression in the definition of tasks. If the specified stream contains no change data, the task skips the current run. And this can not be correct because of the charemski comment.
upvoted 2 times

ivanjs44 2 years, 5 months ago

I think this is also true but the question says: with the LEAST amount of codingSorry, I think STREAM_HAS_DATA can not change the scheduling of a task: This function is intended to be used in the WHEN expression in the definition of tasks. If the specified stream contains no change data, the task skips the current run. And this can not be correct because of the charemski comment.

upvoted 1 times

ivanjs44 2 years, 5 months ago

I think this is also true but the question says: with the LEAST amount of coding

upvoted 1 times

ivanjs44 2 years, 5 months ago

Sorry, I think STREAM_HAS_DATA can not change the scheduling of a task: This function is intended to be used in the WHEN expression in the definition of tasks. If the specified stream contains no change data, the task skips the current run. And this can not be correct because of the charemski comment.

upvoted 1 times

charemski 3 years, 3 months ago

Minimum schedule time for a task is 60 seconds, so in this case only Snowpipe or Materialized View on External Table, in either case using event notification service https://docs.snowflake.com/en/user-guide/tables-external-intro#materialized-views-over-external-tables

upvoted 3 times

hillcat111 3 years, 3 months ago

Answer is A,B and is validatedCOPY command does not exist (option b).

upvoted 1 times

hillcat111 3 years, 3 months ago

Answer is A,B and is validated

upvoted 1 times

dormin 3 years, 2 months ago

COPY command does not exist (option b).

upvoted 2 times

starkbi 3 years, 3 months ago

A: Just define one object: the Pipe object with included COPY INTO ... FROM Stage ... command B: Just define one object: the scheduled Task object with included COPY INTO ... FROM Stage ... command

upvoted 1 times

hillcat111 3 years, 4 months ago

I think A & C are the answersA & E is the right answer

upvoted 2 times

hillcat111 3 years, 4 months ago

I think A & C are the answers

upvoted 2 times

hillcat111 3 years, 4 months ago

A & E is the right answer

upvoted 1 times

Question #33 *Topic 1*

When loading data into a table that captures the load time in a column with a default value of either CURRENT_TIME() or CURRENT_TIMESTAMP() what will occur?

A. All rows loaded using a specific COPY statement will have varying timestamps based on when the rows were inserted.

B. Any rows loaded using a specific COPY statement will have varying timestamps based on when the rows were read from the source.

C. Any rows loaded using a specific COPY statement will have varying timestamps based on when the rows were created in the source.

D. All rows loaded using a specific COPY statement will have the same timestamp value.

Suggested Answer: *C*

Community vote distribution

D (100%)

ArunH Most Recent 10 months, 1 week ago

It gives the timestamp when the Select inside the COPY is executed

upvoted 1 times

charemski 2 years, 3 months ago

Validated and a bit helpful: https://docs.snowflake.com/en/user-guide/data-load-transform#current-time-current-timestamp-default-column-values

upvoted 2 times

hillcat111 2 years, 3 months ago

Answer is D and is validated

upvoted 2 times

hillcat111 2 years, 4 months ago

I think C is the AnswerAnswer is D

upvoted 1 times

hillcat111 2 years, 4 months ago

I think C is the Answer

upvoted 1 times

hillcat111 2 years, 4 months ago

Answer is D

upvoted 1 times

Question #34 *Topic 1*

How does a standard virtual warehouse policy work in Snowflake?

A. It conserves credits by keeping running clusters fully loaded rather than starting additional clusters.

B. It starts only if the system estimates that there is a query load that will keep the cluster busy for at least 6 minutes.

C. It starts only f the system estimates that there is a query load that will keep the cluster busy for at least 2 minutes.

D. It prevents or minimizes queuing by starting additional clusters instead of conserving credits.

Suggested Answer: *D*

Community vote distribution

D (100%)

hillcat111 Most Recent 2 years, 3 months ago

Answer is D and is validated

upvoted 1 times

victorleonis 2 years, 4 months ago

Standard (default) Prevents/minimizes queuing by favoring starting additional clusters over conserving credits.

upvoted 2 times

HighBMI 2 years, 4 months ago

https://docs.snowflake.com/en/user-guide/warehouses-multicluster#label-mcw-scaling-policies

upvoted 2 times

Question #35 *Topic 1*

Which feature provides the capability to define an alternate cluster key for a table with an existing cluster key?

A. External table

B. Materialized view

C. Search optimization

D. Result cache

Suggested Answer: *B*

Community vote distribution

B (100%)

ywan1600 Most Recent 1 year, 7 months ago

The answer is B

upvoted 1 times

hillcat111 2 years, 3 months ago

Answer is B and is validated

upvoted 2 times

starkbi 2 years, 3 months ago

B is the answer

upvoted 1 times

hillcat111 2 years, 4 months ago

i think its B

upvoted 1 times

hillcat111 2 years, 4 months ago

i think its BAnswer is A

upvoted 1 times

hillcat111 2 years, 4 months ago

Answer is A

upvoted 1 times

Question #36 *Topic 1*

An Architect would like to save quarter-end financial results for the previous six years.
Which Snowflake feature can the Architect use to accomplish this?

A. Search optimization service

B. Materialized view

C. Time Travel

D. Zero-copy cloning

E. Secure views

Suggested Answer: *D*

Community vote distribution

D (100%)

Saai123 Most Recent 9 months ago

D is correct(Note materialized view should be avoided as it may change the data as the new data arrives even without the quater end)

upvoted 1 times

hillcat111 2 years, 3 months ago

Answer is D and is validated

upvoted 2 times

starkbi 2 years, 3 months ago

Answer is D. Clone table and afterwards delete data older than the previous 6 years.

upvoted 2 times

hillcat111 2 years, 4 months ago

I think its DChanging my answer to C time travel

upvoted 1 times

hillcat111 2 years, 4 months ago

I think its D

upvoted 1 times

hillcat111 2 years, 4 months ago

Changing my answer to C time travel

upvoted 1 times

Question #37 *Topic 1*

A company is using a Snowflake account in Azure. The account has SAML SSO set up using ADFS as a SCIM identity provider. To validate Private Link connectivity, an Architect performed the following steps:
Confirmed Private Link URLs are working by logging in with a username/password account
Verified DNS resolution by running nslookups against Private Link URLs
Validated connectivity using SnowCD
Disabled public access using a network policy set to use the company's IP address range
However, the following error message is received when using SSO to log into the company account:
IP XX.XXX.XX.XX is not allowed to access snowflake. Contact your local security administrator.
What steps should the Architect take to resolve this error and ensure that the account is accessed using only Private Link? (Choose two.)

A. Alter the Azure security integration to use the Private Link URLs.

B. Add the IP address in the error message to the allowed list in the network policy.

C. Generate a new SCIM access token using system$generate_scim_access_token and save it to Azure AD.

D. Update the configuration of the Azure AD SSO to use the Private Link URLs.

E. Open a case with Snowflake Support to authorize the Private Link URLs' access to the account.

Suggested Answer: *B, C*

Community vote distribution

AD (71%) BD (29%)

785de4c Most Recent 1 year, 2 months ago
https://docs.snowflake.com/en/user-guide/privatelink-azure#configure-access-to-snowflake-with-azure-private-link https://docs.snowflake.com/en/user-guide/admin-security-fed-auth-overview#label-sso-private-connectivity https://docs.snowflake.com/en/user-guide/admin-security-fed-auth-security-integration#create-a-saml2-security-integration
upvoted 1 times

laksnarn 1 year, 8 months ago
After changing the private links , if the SSO want to work the first step is update the security integration and other mandatory step is update the config in azure ad level sso.
upvoted 4 times

jjordan 2 years, 3 months ago
Selected Answer: BD
B&D is correct.
upvoted 2 times

jjordan 2 years, 3 months ago
B&D is correct.Blocked list gets validated first, since the pub access blocked, doesn't matter if anything is in the allowed list
upvoted 2 times

potatosack 1 year, 8 months ago
Blocked list gets validated first, since the pub access blocked, doesn't matter if anything is in the allowed list
upvoted 1 times

hillcat111 2 years, 9 months ago
Answer is B,D and is validated
upvoted 1 times

starkbi 2 years, 9 months ago
I assume B & D
upvoted 1 times

hillcat111 2 years, 10 months ago
i think its B & DChanging answer to A&D
upvoted 1 times

hillcat111 2 years, 10 months ago

i think its B & D

upvoted 1 times

hillcat111 2 years, 10 months ago

Changing answer to A&D

upvoted 2 times

Question #38 *Topic 1*

Which steps are recommended best practices for prioritizing cluster keys in Snowflake? (Choose two.)

A. Choose columns that are frequently used in join predicates.

B. Choose lower cardinality columns to support clustering keys and cost effectiveness.

C. Choose TIMESTAMP columns with nanoseconds for the highest number of unique rows.

D. Choose cluster columns that are most actively used in selective filters.

E. Choose cluster columns that are actively used in the GROUP BY clauses.

Suggested Answer: *A, D*

Community vote distribution

AD (100%)

hillcat111 Most Recent 2 years, 3 months ago

Answer is A,D and is validated

upvoted 1 times

callipso21 2 years, 5 months ago

https://docs.snowflake.com/en/user-guide/tables-clustering-keys 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

Question #39 *Topic 1*

Which Snowflake data modeling approach is designed for BI queries?

A. 3 NF

B. Star schema

C. Data Vault

D. Snowflake schema

Suggested Answer: *D*

Community vote distribution

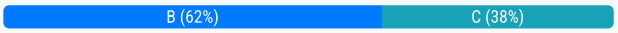

785de4c Most Recent 1 year, 2 months ago

In data models for general BI use, it is important to minimize the number of joins. The star schema has a simple structure with a single fact table at the center, directly connected to multiple denormalized dimension tables. This design requires only one level of joins between the fact table and each dimension, resulting in fewer joins in queries and therefore better performance. https://www.databricks.com/glossary/star-schema

upvoted 1 times

jjordan 2 years, 3 months ago

B is correct.

upvoted 4 times

BDaly 2 years, 4 months ago

B. Star schema The star schema is a common data modeling approach used specifically for optimizing BI queries. It features a central fact table linked to dimension tables, which radiate from the fact like points on a star. This design simplifies queries, improves data retrieval speed, and is highly effective for BI applications where performance and simplicity are key. While Snowflake's architecture can support various data modeling techniques like 3NF (Third Normal Form), Data Vault, and Snowflake schema (which is a variation of star schema with normalized dimension tables), the star schema remains the preferred choice for BI-focused workloads due to its straightforward structure and query efficiency.

upvoted 1 times

hillcat111 2 years, 9 months ago

Answer is B and is validated

upvoted 3 times

starkbi 2 years, 9 months ago

https://www.snowflake.com/blog/support-multiple-data-modeling-approaches-with-snowflake/

upvoted 2 times

hillcat111 2 years, 10 months ago

I think its B

upvoted 1 times

victorleonis 2 years, 10 months ago

I guess that is C https://community.snowflake.com/s/article/Snowflake-Data-Vault-Reference-Architecture

upvoted 1 times

Alachramkowa 2 years, 11 months ago

isn't star schema a better option?

upvoted 2 times

Question #40 *Topic 1*

How is the change of local time due to daylight savings time handled in Snowflake tasks? (Choose two.)

A. A task scheduled in a UTC-based schedule will have no issues with the time changes.

B. Task schedules can be designed to follow specified or local time zones to accommodate the time changes.

C. A task will move to a suspended state during the daylight savings time change.

D. A frequent task execution schedule like minutes may not cause a problem, but will affect the task history.

E. A task schedule will follow only the specified time and will fail to handle lost or duplicated hours.

Suggested Answer: *A, B*

Community vote distribution



starkbi Highly Voted 3 years, 3 months ago

Against B: "A scheduled task runs according to the specified cron expression in the local time for a given time zone. Special care should be taken with regard to scheduling tasks for time zones that recognize daylight saving time. Tasks scheduled during specific times on days when the transition from standard time to daylight saving time (or the reverse) occurs can have unexpected behaviors."

upvoted 7 times

laksnarn Most Recent 2 years, 3 months ago

These options are correct because: Using a UTC-based schedule (Option A) avoids issues with daylight saving time changes as UTC does not observe daylight saving time. Snowflake allows task schedules to follow specified or local time zones (Option B), which can accommodate daylight saving time changes by adjusting the schedule accordingly.

upvoted 1 times

hillcat111 3 years, 3 months ago

Answer is A,E and is validated

upvoted 2 times

victorleonis 3 years, 4 months ago

I guess that is A, B

upvoted 3 times

Alachramkowa 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/tasks-intro#task-scheduling-and-daylight-saving-time AE

upvoted 4 times

Alachramkowa 3 years, 5 months ago

https://docs.snowflake.com/en/user-guide/tasks-intro#task-scheduling-and-daylight-saving-time AEThis example provided in your link is below: "During the autumn change from daylight saving time to standard time, a task scheduled to start at 1 AM in the America/Los_Angeles time zone (i.e. 0 1 * * * America/Los_Angeles) would run twice: once at 1 AM and then again when 1:59:59 AM shifts to 1:00:00 AM local time. That is, there are two points in time when the local time is 1 AM." option E states "fail to handle" it is not failing to handle the task as it is executing it twice.

upvoted 4 times

h8er 2 years, 5 months ago

This example provided in your link is below: "During the autumn change from daylight saving time to standard time, a task scheduled to start at 1 AM in the America/Los_Angeles time zone (i.e. 0 1 * * * America/Los_Angeles) would run twice: once at 1 AM and then again when 1:59:59 AM shifts to 1:00:00 AM local time. That is, there are two points in time when the local time is 1 AM." option E states "fail to handle" it is not failing to handle the task as it is executing it twice.

upvoted 1 times

Question #41 *Topic 1*

An Architect needs to grant a group of ORDER_ADMIN users the ability to clean old data in an ORDERS table (deleting all records older than 5 years), without granting any privileges on the table. The group's manager (ORDER_MANAGER) has full DELETE privileges on the table.
How can the ORDER_ADMIN role be enabled to perform this data cleanup, without needing the DELETE privilege held by the ORDER_MANAGER role?

A. Create a stored procedure that runs with caller's rights, including the appropriate "> 5 years" business logic, and grant USAGE on this procedure to ORDER_ADMIN. The ORDER_MANAGER role owns the procedure.

B. Create a stored procedure that can be run using both caller's and owner's rights (allowing the user to specify which rights are used during execution), and grant USAGE on this procedure to ORDER_ADMIN. The ORDER_MANAGER role owns the procedure.

C. Create a stored procedure that runs with owner's rights, including the appropriate "> 5 years" business logic, and grant USAGE on this procedure to ORDER_ADMIN. The ORDER_MANAGER role owns the procedure.

D. This scenario would actually not be possible in Snowflake – any user performing a DELETE on a table requires the DELETE privilege to be granted to the role they are using.

Suggested Answer: *C*

Community vote distribution

C (100%)

MSIDDIQUI18 Most Recent 1 year, 11 months ago

Answer is "C", You want to delegate a task(s) to another user(s) who will run with the owner's privileges, not the caller's own privileges. For example, if you want a user without DELETE privilege on a table to be able to call a stored procedure that deletes old data, but not current data, then you probably want to use an owner's rights stored procedure. That procedure will contain a DELETE statement that includes a filter (a WHERE clause) to control which data can be deleted through the filter. If you need an owner's rights stored procedure to perform actions on a table, view, or function that the caller has the privileges to access, you can have the caller pass a reference to that table, view, or function.

upvoted 1 times

hillcat111 3 years, 3 months ago

Answer is C and is validated

upvoted 2 times

serg_khar 3 years, 5 months ago

https://docs.snowflake.com/en/sql-reference/stored-procedures-rights#choosing-between-owner-s-rights-and-caller-s-rights

upvoted 3 times

callipso21 3 years, 5 months ago

Tested and C is the correct answer

upvoted 3 times

Question #42 *Topic 1*

A company's daily Snowflake workload consists of a huge number of concurrent queries triggered between 9pm and 11pm. At the individual level, these queries are smaller statements that get completed within a short time period.
What configuration can the company's Architect implement to enhance the performance of this workload? (Choose two.)

A. Enable a multi-clustered virtual warehouse in maximized mode during the workload duration.

B. Set the MAX_CONCURRENCY_LEVEL to a higher value than its default value of 8 at the virtual warehouse level.

C. Increase the size of the virtual warehouse to size X-Large.

D. Reduce the amount of data that is being processed through this workload.

E. Set the connection timeout to a higher value than its default.

Suggested Answer: *A, C*

Community vote distribution

AB (100%)

Arkitekt Most Recent 1 year, 5 months ago
It's all about concurrency of many small queries, and that's the point of a multi-clustered warehouse. Warehouse size is more about handling large single queries.
upvoted 1 times

mrango 2 years, 1 month ago
A,B are correct
upvoted 3 times

hillcat111 2 years, 3 months ago
Answer is A,B and is validated
upvoted 1 times

hillcat111 2 years, 4 months ago
correct A&CChanging to A&D
upvoted 1 times

hillcat111 2 years, 4 months ago
correct A&C
upvoted 1 times

hillcat111 2 years, 4 months ago
Changing to A&D
upvoted 1 times

Question #43 *Topic 1*

Which system functions does Snowflake provide to monitor clustering information within a table (Choose two.)

A. SYSTEM$CLUSTERING_INFORMATION

B. SYSTEM$CLUSTERING_USAGE

C. SYSTEM$CLUSTERING_DEPTH

D. SYSTEM$CLUSTERING_KEYS

E. SYSTEM$CLUSTERING_PERCENT

Suggested Answer: *A, C*

Community vote distribution

AC (100%)

hillcat111 Most Recent 2 years, 3 months ago

Answer is A,C and is validated

upvoted 1 times

HighBMI 2 years, 4 months ago

https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions

upvoted 1 times

Question #44 *Topic 1*

A company has a table with that has corrupted data, named Data. The company wants to recover the data as it was 5 minutes ago using cloning and Time Travel.
What command will accomplish this?

A. CREATE CLONE TABLE Recover_Data FROM Data AT(OFFSET => -60*5);

B. CREATE CLONE Recover_Data FROM Data AT(OFFSET => -60*5);

C. CREATE TABLE Recover_Data CLONE Data AT(OFFSET => -60*5);

D. CREATE TABLE Recover Data CLONE Data AT(TIME => -60*5);

Suggested Answer: *C*

Community vote distribution

C (100%)

hillcat111 Most Recent 2 years, 3 months ago

Answer is C and is validated

upvoted 1 times

callipso21 2 years, 5 months ago

https://docs.snowflake.com/en/user-guide/data-time-travel Section - Cloning Historical Objects

upvoted 2 times

Question #45 *Topic 1*

A company has an inbound share set up with eight tables and five secure views. The company plans to make the share part of its production data pipelines.
Which actions can the company take with the inbound share? (Choose two.)

A. Clone a table from a share.

B. Grant modify permissions on the share.

C. Create a table from the shared database.

D. Create additional views inside the shared database.

E. Create a table stream on the shared table.

Suggested Answer: *C, E*

Community vote distribution

CE (100%)

MSIDDIQUI18 Most Recent 1 year, 11 months ago

Correct Answer: CE CREATE DATABASE c1 FROM SHARE provider1.share1; Creating streams on shared objects (secure views or tables) enables you to track data manipulation language (DML) changes made in those objects. CREATE STREAM ON TABLE . . ;

upvoted 1 times

Pk2287 2 years, 2 months ago

Answer is C & E.. A,B and D options are not possible considering share data is read only

upvoted 1 times

drunk_goat82 2 years, 5 months ago

CAN'T CLONE A TABLE from a share.

upvoted 1 times

DataKelpie 3 years, 2 months ago

On the basis of this documentation, the answer is C and E https://docs.snowflake.com/en/user-guide/data-share-consumers#step-1-create-a-database-from-the-share The reasoning is as follows: Cloning of shared databases is not supported They cannot be modified Streaming (CDC) of updates is supported

upvoted 2 times

hillcat111 3 years, 4 months ago

Answer is A&C

upvoted 1 times

Question #46 *Topic 1*

A company has several sites in different regions from which the company wants to ingest data.
Which of the following will enable this type of data ingestion?

A. The company must have a Snowflake account in each cloud region to be able to ingest data to that account.

B. The company must replicate data between Snowflake accounts.

C. The company should provision a reader account to each site and ingest the data through the reader accounts.

D. The company should use a storage integration for the external stage.

Suggested Answer: *A*

Community vote distribution

D (100%)

jjordan Most Recent 2 years, 3 months ago

D is correct.

upvoted 3 times

DataKelpie 2 years, 8 months ago

D would seem to be the correct answer because the question is about data that has not already been ingested into Snowflake. Consequently cross region replication is not relevant.

upvoted 3 times

hillcat111 2 years, 9 months ago

Answer is D and is validated

upvoted 4 times

starkbi 2 years, 9 months ago

I assume B

upvoted 1 times

hillcat111 2 years, 10 months ago

changing to D

upvoted 2 times

hillcat111 2 years, 10 months ago

Answer is C

upvoted 1 times

Question #47 *Topic 1*

What Snowflake features should be leveraged when modeling using Data Vault?

A. Snowflake's support of multi-table inserts into the data model's Data Vault tables

B. Data needs to be pre-partitioned to obtain a superior data access performance

C. Scaling up the virtual warehouses will support parallel processing of new source loads

D. Snowflake's ability to hash keys so that hash key joins can run faster than integer joins

Suggested Answer: *C*

Community vote distribution

A (100%)

1940f9d Most Recent 1 year, 1 month ago

It should be A. it exhibits behavior of data loading for data vault

upvoted 1 times

chiiiik 2 years ago

Answer is A, C is wrong cause scaling up improves query performance, it should be scaling out as data vault require high concurrent query execution

upvoted 2 times

hillcat111 2 years, 9 months ago

Answer is A and is validated

upvoted 2 times

hillcat111 2 years, 10 months ago

changing to A

upvoted 1 times

hillcat111 2 years, 10 months ago

this seems like the correct answer

upvoted 1 times

Question #48 *Topic 1*

A company's client application supports multiple authentication methods, and is using Okta.
What is the best practice recommendation for the order of priority when applications authenticate to Snowflake?

A. 1. OAuth (either Snowflake OAuth or External OAuth) 2. External browser 3. Okta native authentication 4. Key Pair Authentication, mostly used for service account users 5. Password

B. 1. External browser, SSO 2. Key Pair Authentication, mostly used for development environment users 3. Okta native authentication 4. OAuth (ether Snowflake OAuth or External OAuth) 5. Password

C. 1. Okta native authentication 2. Key Pair Authentication, mostly used for production environment users 3. Password 4. OAuth (either Snowflake OAuth or External OAuth) 5. External browser, SSO

D. 1. Password 2. Key Pair Authentication, mostly used for production environment users 3. Okta native authentication 4. OAuth (either Snowflake OAuth or External OAuth) 5. External browser, SSO

Suggested Answer: *A*

Community vote distribution

A (100%)

Fabrisa Highly Voted 2 years, 5 months ago
I would choose A https://community.snowflake.com/s/article/Snowflake-Security-Overview-and-Best-Practices
upvoted 6 times

hillcat111 Most Recent 2 years, 3 months ago
Answer is A and is validated
upvoted 2 times

Fnascimento 2 years, 4 months ago
Is option A
upvoted 3 times

Question #49 *Topic 1*

What is a valid object hierarchy when building a Snowflake environment?

A. Account --> Database --> Schema --> Warehouse

B. Organization --> Account --> Database --> Schema --> Stage

C. Account --> Schema > Table --> Stage

D. Organization --> Account --> Stage --> Table --> View

Suggested Answer: *B*

Community vote distribution

B (100%)

ywan1600 Most Recent 1 year, 7 months ago

Answer is B

upvoted 1 times

hillcat111 2 years, 3 months ago

Answer is B and is validated

upvoted 3 times

hillcat111 2 years, 4 months ago

i would choose A

upvoted 1 times

hillcat111 2 years, 4 months ago

i would choose AChanging to B

upvoted 1 times

hillcat111 2 years, 3 months ago

Changing to B

upvoted 1 times

Question #50 *Topic 1*

Which of the following are characteristics of Snowflake's parameter hierarchy?

A. Session parameters override virtual warehouse parameters.

B. Virtual warehouse parameters override user parameters.

C. Table parameters override virtual warehouse parameters.

D. Schema parameters override account parameters.

Suggested Answer: *D*

Community vote distribution

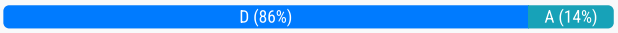

charemski Highly Voted 2 years, 3 months ago

Session parameters not always override Virtual Warehouse parameters, see i.e. here: https://docs.snowflake.com/sql-reference/parameters#statement-timeout-in-seconds Only account->database->schema->table always keeps hierarchy of parameters

upvoted 6 times

anikl Most Recent 1 year, 11 months ago

When the parameter like STATEMENT_QUEUED_TIMEOUT_IN_SECONDS is set for both a warehouse and a session, the lowest non-zero value is enforced therefore A is not correct. D is my answer

upvoted 2 times

hillcat111 2 years, 3 months ago

Answer is A and is validated

upvoted 2 times

callipso21 2 years, 5 months ago

https://docs.snowflake.com/en/user-guide/admin-account-management Snowflake provides three types of parameters that can be set for your account: - Account parameters that affect your entire account. - Session parameters that default to users and their sessions. - Object parameters that default to objects (warehouses, databases, schemas, and tables).

upvoted 1 times

Question #51 *Topic 1*

A Snowflake Architect is designing a multi-tenant application strategy for an organization in the Snowflake Data Cloud and is considering using an Account Per Tenant strategy.
Which requirements will be addressed with this approach? (Choose two.)

A. There needs to be fewer objects per tenant.

B. Security and Role-Based Access Control (RBAC) policies must be simple to configure.

C. Compute costs must be optimized.

D. Tenant data shape may be unique per tenant.

E. Storage costs must be optimized.

Suggested Answer: *C, E*

Community vote distribution

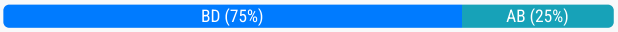

Fabrisa Highly Voted 3 years, 5 months ago
I would choose B and D... https://developers.snowflake.com/wp-content/uploads/2021/05/Design-Patterns-for-Building-Multi-Tenant-Applications-on-Snowflake.pdf
upvoted 9 times

Fabrisa Highly Voted 3 years, 5 months ago
I would choose B and D... https://developers.snowflake.com/wp-content/uploads/2021/05/Design-Patterns-for-Building-Multi-Tenant-Applications-on-Snowflake.pdfI agree here.
upvoted 10 times

Corleone433 3 years ago
I agree here.
upvoted 1 times

blink291 1 year, 10 months ago
I would choose A, B. Option D can also be acheived using OPT as mentioned here (Page 3 https://developers.snowflake.com/wp-content/uploads/2021/05/Design-Patterns-for-Building-Multi-Tenant-Applications-on-Snowflake.pdf)
upvoted 1 times

chiiiik 2 years, 6 months ago
I think we can exclude C & E cause they are not requirements well solved by using APT, I'm not sure about A, but it seems unnecessary to use APT with few objects. Potential is A cause RBAC is not required to isolate tenant (Refer Part Three: Evaluating Tenancy Models in https://developers.snowflake.com/wp-content/uploads/2021/05/Design-Patterns-for-Building-Multi-Tenant-Applications-on-Snowflake.pdf). And the left is D, I think it's reasonable to use APT
upvoted 3 times

Atomic_Gecko 2 years, 7 months ago
C & D. Compute costs are to be considered when using multiple accounts, D for data shape being different
upvoted 1 times

hillcat111 3 years, 3 months ago
Answer is C,E and is validatedWhy would storage costs be optimized?!
upvoted 1 times

hillcat111 3 years, 3 months ago
Answer is C,E and is validated
upvoted 1 times

ukpino Most Recent 1 year, 9 months ago
Why would storage costs be optimized?!
upvoted 1 times

hillcat111 3 years, 3 months ago
changing to B&D
upvoted 2 times

hillcat111 3 years, 4 months ago

I think its A&D

upvoted 1 times

Question #52 *Topic 1*

An Architect has been asked to clone schema STAGING as it looked one week ago, Tuesday June 1st at 8:00 AM, to recover some objects.
The STAGING schema has 50 days of retention.
The Architect runs the following statement:
CREATE SCHEMA STAGING_CLONE CLONE STAGING at (timestamp => '2021-06-01 08:00:00');
The Architect receives the following error: Time travel data is not available for schema STAGING. The requested time is either beyond the allowed time travel period or before the object creation time.
The Architect then checks the schema history and sees the following:

CREATED_ON|NAME|DROPPED_ON -
2021-06-02 23:00:00 | STAGING | NULL
2021-05-01 10:00:00 | STAGING | 2021-06-02 23:00:00
How can cloning the STAGING schema be achieved?

A. Undrop the STAGING schema and then rerun the CLONE statement.

B. Modify the statement: CREATE SCHEMA STAGING_CLONE CLONE STAGING at (timestamp => '2021-05-01 10:00:00');

C. Rename the STAGING schema and perform an UNDROP to retrieve the previous STAGING schema version, then run the CLONE statement.

D. Cloning cannot be accomplished because the STAGING schema version was not active during the proposed Time Travel time period.

Suggested Answer: *C*

Community vote distribution

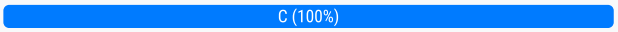

ywan1600 Most Recent 1 year, 7 months ago
Answer is C
upvoted 1 times

hillcat111 2 years, 3 months ago
Answer is C and is validated
upvoted 1 times

hillcat111 2 years, 4 months ago
Answer is CChanging answer to D
upvoted 1 times

hillcat111 2 years, 4 months ago
Answer is C
upvoted 1 times

hillcat111 2 years, 4 months ago
Changing answer to D
upvoted 1 times

Question #53 *Topic 1*

What are purposes for creating a storage integration? (Choose three.)

A. Control access to Snowflake data using a master encryption key that is maintained in the cloud provider's key management service.

B. Store a generated identity and access management (IAM) entity for an external cloud provider regardless of the cloud provider that hosts the Snowflake account.

C. Support multiple external stages using one single Snowflake object.

D. Avoid supplying credentials when creating a stage or when loading or unloading data.

E. Create private VPC endpoints that allow direct, secure connectivity between VPCs without traversing the public internet.

F. Manage credentials from multiple cloud providers in one single Snowflake object.

Suggested Answer: *B, C, D*

Community vote distribution

BCD (100%)

ywan1600 Most Recent 1 year, 7 months ago

Answer is B,C,D

upvoted 1 times

OscarWest2023 2 years, 1 month ago

B,C,D https://docs.snowflake.com/en/sql-reference/sql/create-storage-integration 1 - A storage integration is a Snowflake object that stores a generated identity and access management (IAM) entity for your external cloud storage. 2 - This option allows users to avoid supplying credentials when creating stages or when loading or unloading data. 3 - A single storage integration can support multiple external stages. The URL in the stage definition must align with the storage location specified for the STORAGE_ALLOWED_LOCATIONS parameter.

upvoted 2 times

hillcat111 2 years, 9 months ago

Answer is B,C,D and is validated

upvoted 4 times

hillcat111 2 years, 10 months ago

changing to C,D,F

upvoted 1 times

hillcat111 2 years, 10 months ago

Answer here B, D, F

upvoted 1 times

Question #54 *Topic 1*

A healthcare company is deploying a Snowflake account that may include Personal Health Information (PHI). The company must ensure compliance with all relevant privacy standards.
Which best practice recommendations will meet data protection and compliance requirements? (Choose three.)

A. Use, at minimum, the Business Critical edition of Snowflake.

B. Create Dynamic Data Masking policies and apply them to columns that contain PHI.

C. Use the Internal Tokenization feature to obfuscate sensitive data.

D. Use the External Tokenization feature to obfuscate sensitive data.

E. Rewrite SQL queries to eliminate projections of PHI data based on current_role().

F. Avoid sharing data with partner organizations.

Suggested Answer: *A, B, D*

Community vote distribution

ABD (100%)

ywan1600 Most Recent 1 year, 7 months ago
Answer is A,B,D
upvoted 1 times

OscarWest2023 2 years, 1 month ago
A,B,D. You need External Tokenization and Dynamic data masking to meet the compliance and privacy requirements. While bot of those features are available within Enterprise Edition, Snowflakes documentation specifies to use Business critical where you have PHI and other privacy requirements. Business Critical Edition, 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. It includes all the features and services of Enterprise Edition, with the addition of enhanced security and data protection. In addition, database failover/failback adds support for business continuity and disaster recovery.
https://docs.snowflake.com/en/user-guide/intro-editions
upvoted 2 times

hillcat111 2 years, 9 months ago
Answer is A,B,D and is validated
upvoted 1 times

hillcat111 2 years, 10 months ago
changing to B, C, E
upvoted 1 times

hillcat111 2 years, 10 months ago
Answer B, C & F
upvoted 1 times

Question #55 *Topic 1*

There are two databases in an account, named fin_db and hr_db which contain payroll and employee data, respectively. Accountants and Analysts in the company require different permissions on the objects in these databases to perform their jobs. Accountants need read-write access to fin_db but only require read-only access to hr_db because the database is maintained by human resources personnel.
An Architect needs to create a read-only role for certain employees working in the human resources department.
Which permission sets must be granted to this role?

A. USAGE on database hr_db, USAGE on all schemas in database hr_db, SELECT on all tables in database hr_db

B. USAGE on database hr_db, SELECT on all schemas in database hr_db, SELECT on all tables in database hr_db

C. MODIFY on database hr_db, USAGE on all schemas in database hr_db, USAGE on all tables in database hr_db

D. USAGE on database hr_db, USAGE on all schemas in database hr_db, REFERENCES on all tables in database hr_db

Suggested Answer: *A*

Community vote distribution

A (100%)

ywan1600 Most Recent 1 year, 7 months ago

Answer is A

upvoted 1 times

hillcat111 2 years, 3 months ago

Answer is A and is validated

upvoted 2 times

hillcat111 2 years, 4 months ago

Answer is correct A

upvoted 1 times

Question #56 *Topic 1*

An Architect runs the following SQL query:

How can this query be interpreted?

```
SELECT
   METADATA$FILENAME,
   METADATA$FILE_ROW_NUMBER
FROM @FILEROWS/Food_Reviews.csv
       (file_format=CSV_N)
```

A. FILEROWS is a stage. FILE_ROW_NUMBER is line number in file.

B. FILEROWS is the table. FILE_ROW_NUMBER is the line number in the table.

C. FILEROWS is a file. FILE_ROW_NUMBER is the file format location.

D. FILERONS is the file format location. FILE_ROW_NUMBER is a stage.

Suggested Answer: *A*

Community vote distribution

A (100%)

hillcat111 Most Recent 2 years, 3 months ago

Answer is A and is validated

upvoted 2 times

callipso21 2 years, 5 months ago

METADATA$FILE_ROW_NUMBER Row number for each record in the staged data file.

upvoted 2 times

Question #57 *Topic 1*

An Architect entered the following commands in sequence:

USER1 cannot find the table.

Which of the following commands does the Architect need to run for USER1 to find the tables using the Principle of Least Privilege? (Choose two.)

```
CREATE DATABASE SANDBOX;
CREATE ROLE INTERN;
CREATE TABLE SANDBOX.PUBLIC.AGENDA (ID INT, ITEMS STRING);
GRANT SELECT ON ALL TABLES IN SCHEMA SANDBOX.PUBLIC TO ROLE INTERN;
GRANT ROLE INTERN TO USER USER1;
```

A. GRANT ROLE PUBLIC TO ROLE INTERN;

B. GRANT USAGE ON DATABASE SANDBOX TO ROLE INTERN;

C. GRANT USAGE ON SCHEMA SANDBOX.PUBLIC TO ROLE INTERN;

D. GRANT OWNERSHIP ON DATABASE SANDBOX TO USER INTERN;

E. GRANT ALL PRIVILEGES ON DATABASE SANDBOX TO ROLE INTERN;

Suggested Answer: *B, C*

Community vote distribution



Soumak Most Recent 1 year, 8 months ago

validated

upvoted 1 times

hillcat111 2 years, 3 months ago

Answer is B,C and is validated

upvoted 2 times

hillcat111 2 years, 4 months ago

Answer is correct B & C

upvoted 2 times

Question #58 *Topic 1*

A DevOps team has a requirement for recovery of staging tables used in a complex set of data pipelines. The staging tables are all located in the same staging schema. One of the requirements is to have online recovery of data on a rolling 7-day basis.
After setting up the DATA_RETENTION_TIME_IN_DAYS at the database level, certain tables remain unrecoverable past 1 day.
What would cause this to occur? (Choose two.)

A. The staging schema has not been setup for MANAGED ACCESS.

B. The DATA_RETENTION_TIME_IN_DAYS for the staging schema has been set to 1 day.

C. The tables exceed the 1 TB limit for data recovery.

D. The staging tables are of the TRANSIENT type.

E. The DevOps role should be granted ALLOW_RECOVERY privilege on the staging schema.

Suggested Answer: *B, D*

hillcat111 Most Recent 2 years, 3 months ago
Answer is B,D and is validated
upvoted 3 times

hillcat111 2 years, 4 months ago
Answer seems correct B & D
upvoted 2 times

hillcat111 2 years, 4 months ago
Answer seems correct B & DChanging answer to A & B
upvoted 2 times

hillcat111 2 years, 4 months ago
Changing answer to A & B
upvoted 1 times

Question #59 *Topic 1*

Consider the following COPY command which is loading data with CSV format into a Snowflake table from an internal stage through a data transformation query.

This command results in the following error:
SQL compilation error: invalid parameter 'validation_mode'
Assuming the syntax is correct, what is the cause of this error?

```
copy into home_sales(city, zip, sale_date, price)
from (select t.$1, t.$2, t.$6, t.$7 from @mystage/sales.csv.qz t)
file_format =
(
format_name = mycsvformat
empty_field_as_null = true
field_optionally_enclosed_by = '"'
)
validation_mode - return_all_errors
;
```

A. The VALIDATION_MODE parameter supports COPY statements that load data from external stages only.

B. The VALIDATION_MODE parameter does not support COPY statements with CSV file formats.

C. The VALIDATION_MODE parameter does not support COPY statements that transform data during a load.

D. The value return_all_errors of the option VALIDATION_MODE is causing a compilation error.

Suggested Answer: *C*

Community vote distribution

C (100%)

charemski Highly Voted 3 years, 3 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 5 times

NachoPrendes Most Recent 2 years, 2 months ago

https://docs.snowflake.com/en/sql-reference/sql/copy-into-table#:~:text=Note-,VALIDATION_MODE,-does%20not%20support

upvoted 1 times

hillcat111 3 years, 3 months ago

Answer is C and is validated

upvoted 2 times

hillcat111 3 years, 3 months ago

Fabrisa Is ccrrect C

upvoted 1 times

hillcat111 3 years, 4 months ago

i would choose D validation mode is causing the compiling error

upvoted 1 times

Fabrisa 3 years, 5 months ago

I would choose C... https://docs.snowflake.com/en/sql-reference/sql/copy-into-table#optional-parameters

upvoted 1 times

Fabrisa 3 years, 5 months ago

I would choose C... https://docs.snowflake.com/en/sql-reference/sql/copy-into-table#optional-parametersBased on the link above it's the opposite so "D" is correct.

upvoted 1 times

Katarzyna04 3 years, 3 months ago

Based on the link above it's the opposite so "D" is correct.

upvoted 1 times

Question #60 *Topic 1*

A company is storing large numbers of small JSON files (ranging from 1-4 bytes) that are received from IoT devices and sent to a cloud provider. In any given hour, 100,000 files are added to the cloud provider.
What is the MOST cost-effective way to bring this data into a Snowflake table?

A. An external table

B. A pipe

C. A stream

D. A copy command at regular intervals

Suggested Answer: *B*

Community vote distribution

D (100%)

avijit1982 Most Recent 1 year, 1 month ago
If it is B, then it would not be cost effective
upvoted 1 times

potatosack 1 year, 8 months ago
If only looking at cost-effective, D. If also take into account real-time capability, then it would be B. Pipe can get very expensive handling a lot of small files, its best for real-time/ near real-time scenarios
upvoted 2 times

hillcat111 2 years, 3 months ago
Answer is B and is validated
upvoted 1 times

hillcat111 2 years, 4 months ago
Correct answer
upvoted 1 times

Question #61 *Topic 1*

A company has a Snowflake account named ACCOUNTA in AWS us-east-1 region. The company stores its marketing data in a Snowflake database named MARKET_DB. One of the company's business partners has an account named PARTNERB in Azure East US 2 region. For marketing purposes the company has agreed to share the database MARKET_DB with the partner account.
Which of the following steps MUST be performed for the account PARTNERB to consume data from the MARKET_DB database?

A. Create a new account (called AZABC123) in Azure East US 2 region. From account ACCOUNTA create a share of database MARKET_DB, create a new database out of this share locally in AWS us-east-1 region, and replicate this new database to AZABC123 account. Then set up data sharing to the PARTNERB account.

B. From account ACCOUNTA create a share of database MARKET_DB, and create a new database out of this share locally in AWS us-east-1 region. Then make this database the provider and share it with the PARTNERB account.

C. Create a new account (called AZABC123) in Azure East US 2 region. From account ACCOUNTA replicate the database MARKET_DB to AZABC123 and from this account set up the data sharing to the PARTNERB account.

D. Create a share of database MARKET_DB, and create a new database out of this share locally in AWS us-east-1 region. Then replicate this database to the partner's account PARTNERB.

Suggested Answer: *C*

Community vote distribution



ywan1600 Most Recent 1 year, 7 months ago
Answer is C
upvoted 2 times

h8er 1 year, 11 months ago
I vote C
upvoted 1 times

hillcat111 2 years, 9 months ago
Answer is C and is validated
upvoted 2 times

hillcat111 2 years, 10 months ago
Answer is B
upvoted 1 times

Question #62 *Topic 1*

What does a Snowflake Architect need to consider when implementing a Snowflake Connector for Kafka?

A. Every Kafka message is in JSON or Avro format.

B. The default retention time for Kafka topics is 14 days.

C. The Kafka connector supports key pair authentication, OAUTH, and basic authentication (for example, username and password).

D. The Kafka connector will create one table and one pipe to ingest data for each topic. If the connector cannot create the table or the pipe it will result in an exception.

Suggested Answer: *D*

Community vote distribution

A (100%)

ywan1600 Most Recent 1 year, 7 months ago

The answer should be A

upvoted 1 times

MSIDDIQUI18 1 year, 11 months ago

I think answer is "A": Each Kafka message is passed to Snowflake in JSON format or Avro format. Only issue with "D" answer is the last part of statement, does Snowflake raising the Exception, if unable to connector cannot create the table or the pipe.

upvoted 2 times

cui_li 2 years, 6 months ago

Correction: The answer should be D. The connector creates the following objects for each topic: One internal stage to temporarily store data files for each topic. One pipe to ingest the data files for each topic partition. One table for each topic. If the table specified for each topic does not exist, the connector creates i

upvoted 1 times

cui_li 2 years, 6 months ago

Ye should be A. https://docs.snowflake.com/en/user-guide/kafka-connector-overview

upvoted 1 times

Atomic_Gecko 2 years, 7 months ago

Should be A. Default retention time is 7 days, only key/pair auth is supported and the connector will create a pipe, table and stage

upvoted 1 times

Question #63 *Topic 1*

A group of Data Analysts have been granted the role ANALYST_ROLE. They need a Snowflake database where they can create and modify tables, views, and other objects to load with their own data. The Analysts should not have the ability to give other Snowflake users outside of their role access to this data.

How should these requirements be met?

A. Grant ANALYST_ROLE OWNERSHIP on the database, but make sure that ANALYST_ROLE does not have the MANAGE GRANTS privilege on the account.

B. Grant SYSADMIN OWNERSHIP of the database, but grant the create schema privilege on the database to the ANALYST_ROLE.

C. Make every schema in the database a MANAGED ACCESS schema, owned by SYSADMIN, and grant create privileges on each schema to the ANALYST_ROLE for each type of object that needs to be created.

D. Grant ANALYST_ROLE OWNERSHIP on the database, but grant the OWNERSHIP ON FUTURE [object type]s in database privilege to SYSADMIN.

Suggested Answer: *C*

Community vote distribution

C (100%)

ywan1600 Most Recent 1 year, 7 months ago

C is correct

upvoted 1 times

Atomic_Gecko 2 years, 1 month ago

C is correct

upvoted 1 times

Question #64 *Topic 1*

What considerations need to be taken when using database cloning as a tool for data lifecycle management in a development environment? (Choose two.)

A. Any pipes in the source are not cloned.

B. Any pipes in the source referring to internal stages are not cloned.

C. Any pipes in the source referring to external stages are not cloned.

D. The clone inherits all granted privileges of all child objects in the source object, including the database.

E. The clone inherits all granted privileges of all child objects in the source object, excluding the database.

Suggested Answer: *A, C*

Community vote distribution

BE (100%)

Atomic_Gecko Highly Voted 2 years, 1 month ago

Correct answer is B & E, only pipes referencing external stages are cloned, and the clone inherits all grants to child objects but not the database itself

upvoted 6 times

rvalusa Most Recent 7 months, 1 week ago

When you clone a database or schema, Snowflake does not copy the actual data files sitting inside internal stages. Because the newly cloned internal stage will always be completely empty, a cloned pipe would have absolutely no files to process. To prevent confusion, errors, or accidental duplicate loading from the original source, Snowflake simply skips cloning that pipe altogether.

upvoted 1 times

bizolus 1 year, 4 months ago

B, E, correct answers

upvoted 1 times

potatosack 1 year, 8 months ago

B & E, https://docs.snowflake.com/en/user-guide/object-clone

upvoted 2 times

cui_li 1 year, 10 months ago

Answer should be B & E. https://docs.snowflake.com/en/user-guide/object-clone When a database or schema is cloned, any pipes in the source container that reference an internal (i.e. Snowflake) stage are not cloned. However, any pipes that reference an external stage are cloned. 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: Note that the clone of the container itself (database or schema) does not inherit the privileges granted on the source container.

upvoted 4 times

Question #65 *Topic 1*

Which columns can be included in an external table schema? (Choose three.)

A. VALUE

B. METADATA$ROW_ID

C. METADATA$ISUPDATE

D. METADATA$FILENAME

E. METADATA$FILE_ROW_NUMBER

F. METADATA$EXTERNAL_TABLE_PARTITION

Suggested Answer: *A, D, E*

Community vote distribution

ADE (100%)

Arkitekt Most Recent 1 year, 5 months ago

All external tables include the following columns: VALUE A VARIANT type column that represents a single row in the external file. METADATA$FILENAME A pseudocolumn that identifies the name of each staged data file included in the external table, including its path in the stage. METADATA$FILE_ROW_NUMBER A pseudocolumn that shows the row number for each record in a staged data file.

upvoted 1 times

ywan1600 1 year, 7 months ago

A, D & E is correct : https://docs.snowflake.com/en/user-guide/tables-external-intro

upvoted 1 times

Atomic_Gecko 2 years, 1 month ago

A, D & E is correct : https://docs.snowflake.com/en/user-guide/tables-external-intro

upvoted 4 times

Question #66 *Topic 1*

Which SQL ALTER command will MAXIMIZE memory and compute resources for a Snowpark stored procedure when executed on the snowpark_opt_wh warehouse?

A. alter warehouse snowpark_opt_wh set max_concurrency_level = 1;

B. alter warehouse snowpark_opt_wh set max_concurrency_level = 2;

C. alter warehouse snowpark_opt_wh set max_concurrency_level = 8;

D. alter warehouse snowpark_opt_wh set max_concurrency_level = 16;

Suggested Answer: *A*

Community vote distribution

A (100%)

ywan1600 Most Recent 1 year, 7 months ago

A is correct

upvoted 1 times

swe12 2 years, 1 month ago

A is right Set the warehouse MAX_CONCURRENCY_LEVEL parameter to 1 to maximize memory and compute resources for a given Snowpark stored procedure. https://docs.snowflake.com/en/developer-guide/snowpark/python/python-snowpark-training-ml

upvoted 3 times

Question #67 *Topic 1*

An Architect clones a database and all of its objects, including tasks. After the cloning, the tasks stop running.

Why is this occurring?

A. Tasks cannot be cloned.

B. The objects that the tasks reference are not fully qualified.

C. Cloned tasks are suspended by default and must be manually resumed.

D. The Architect has insufficient privileges to alter tasks on the cloned database.

Suggested Answer: *C*

Community vote distribution

C (100%)

ywan1600 Most Recent 1 year, 7 months ago

C is correct

upvoted 1 times

swe12 2 years, 1 month ago

C is right. When a database or schema that contains tasks is cloned, the tasks in the clone are suspended by default.

https://docs.snowflake.com/en/user-guide/object-clone

upvoted 3 times

Question #68 *Topic 1*

Which technique will efficiently ingest and consume semi-structured data for Snowflake data lake workloads?

A. IDEF1X

B. Schema-on-write

C. Schema-on-read

D. Information schema

Suggested Answer: *C*

Community vote distribution

C (100%)

Soumak Most Recent 1 year, 8 months ago

schema on read is relevant to semi-structured data

upvoted 2 times

67422cb 1 year, 10 months ago

schema on read (c) is related to semi-structured data, but it is an odd question dealing with ingesting and consuming. It's not the schema on write will inefficiently ingest it, it just won't be relevant for JSON at all.

upvoted 1 times

Question #69 *Topic 1*

Is it possible for a data provider account with a Snowflake Business Critical edition to share data with an Enterprise edition data consumer account?

A. A Business Critical account cannot be a data sharing provider to an Enterprise consumer. Any consumer accounts must also be Business Critical.

B. If a user in the provider account with role authority to CREATE or ALTER SHARE adds an Enterprise account as a consumer, it can import the share.

C. If a user in the provider account with a share owning role sets SHARE_RESTRICTIONS to False when adding an Enterprise consumer account, it can import the share.

D. If a user in the provider account with a share owning role which also has OVERRIDE SHARE RESTRICTIONS privilege SHARE_RESTRICTIONS set to False when adding an Enterprise consumer account, it can import the share.

Suggested Answer: *D*

Community vote distribution

D (100%)

Arkitekt Most Recent 1 year, 5 months ago

To allow sharing data from a Business Critical to a non-Business Critical account, a user with a role granted the OVERRIDE SHARE RESTRICTIONS global privilege can specify the SHARE_RESTRICTIONS parameter for a specific share offered by their provider account.
https://docs.snowflake.com/en/user-guide/override_share_restrictions

upvoted 1 times

Soumak 1 year, 8 months ago

validated

upvoted 2 times

swe12 2 years ago

D is right https://docs.snowflake.com/en/user-guide/override_share_restrictions

upvoted 3 times

Question #70 *Topic 1*

Which of the following ingestion methods can be used to load near real-time data by using the messaging services provided by a cloud provider?

A. Snowflake Connector for Kafka

B. Snowflake streams

C. Snowpipe

D. Spark

Suggested Answer: *A*

Community vote distribution

C (100%)

Soumak Most Recent 1 year, 8 months ago

snowpipe looks to be the correct answer as it can easily integrate with SQS for streaming data ingestion.

upvoted 2 times

ukpino 1 year, 9 months ago

I would go C. Automating Snowpipe using cloud messaging Automated data loads leverage event notifications for cloud storage to inform Snowpipe of the arrival of new data files to load. Snowpipe polls the event notifications from a queue. https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro

upvoted 2 times

isadd_ 2 years, 2 months ago

A, Kafka is the messaging service so I would use the Snowflake Connector for Kafka

upvoted 1 times

DBSS 2 years, 5 months ago

C Automating Snowpipe using Amazon SQS notifications works well https://docs.snowflake.com/en/user-guide/data-load-snowpipe-auto-s3

upvoted 1 times

Question #71 *Topic 1*

An Architect is designing a file ingestion recovery solution. The project will use an internal named stage for file storage. Currently, in the case of an ingestion failure, the Operations team must manually download the failed file and check for errors.

Which downloading method should the Architect recommend that requires the LEAST amount of operational overhead?

A. Use the Snowflake Connector for Python, connect to remote storage and download the file.

B. Use the GET command in SnowSQL to retrieve the file.

C. Use the GET command in Snowsight to retrieve the file.

D. Use the Snowflake API endpoint and download the file.

Suggested Answer: *B*

Community vote distribution

B (100%)

ywan1600 Most Recent 1 year, 7 months ago

SnowSQL is Snowflake's command-line client. The GET command is specifically designed to download files from Snowflake internal stages to a local file system.

upvoted 1 times

isadd_ 2 years, 2 months ago

B, Snowsight would require to manually download the file each time and the other two options would increase the operational overhead

upvoted 1 times

Question #72 *Topic 1*

A table for IOT devices that measures water usage is created. The table quickly becomes large and contains more than 2 billion rows.

The general query patterns for the table are:

1. DeviceId, IOT_timestamp and CustomerId are frequently used in the filter predicate for the select statement
2. The columns City and DeviceManufacturer are often retrieved
3. There is often a count on UniqueId

Which field(s) should be used for the clustering key?

```
create table water_iot (
UniqueId number,
DeviceId varchar(20),
DeviceManufacturer varchar(50)
CustomerId varchar(20),
IOT_timestamp timestamp_ntz,
City varchar(80),
Location varchar(50)
)
```

A. IOT_timestamp

B. City and DeviceManufacturer

C. DeviceId and CustomerId

D. UniqueId

Suggested Answer: *C*

Community vote distribution

C (100%)

Soumak Most Recent 1 year, 8 months ago

medium cardinal columns, timestamp will have very high cardinality

upvoted 2 times

cui_li 1 year, 10 months ago

Should be C: https://docs.snowflake.com/en/user-guide/tables-clustering-keys

upvoted 2 times

Question #73 *Topic 1*

Which Snowflake objects can be used in a data share? (Choose two.)

A. Standard view

B. Secure view

C. Stored procedure

D. External table

E. Stream

Suggested Answer: *B, D*

Community vote distribution

BD (100%)

ArunH Most Recent 10 months, 1 week ago

B &D are correct as they are most popular. A, a standard view also can be shared with SECURE_OBJECTS_ONLY=FALSE parameter on the share.

upvoted 1 times

ywan1600 1 year, 7 months ago

BD are correct

upvoted 1 times

MSIDDIQUI18 1 year, 10 months ago

Correct Answer: BD Databases, Table, Dynamic table, External table, Iceberg table, Secure Views, Secure Materialized views, Secure user define function (UDFs). All DB objects shared between account are READ ONLY. i.e the objects CANNOT be Modified, or Delete, included adding or Modifying table data.

upvoted 1 times

swe12 2 years, 7 months ago

B D are right https://docs.snowflake.com/en/user-guide/data-sharing-intro

upvoted 3 times

Question #74 *Topic 1*

A company has an external vendor who puts data into Google Cloud Storage. The company's Snowflake account is set up in Azure.

What would be the MOST efficient way to load data from the vendor into Snowflake?

A. Ask the vendor to create a Snowflake account, load the data into Snowflake and create a data share.

B. Create an external stage on Google Cloud Storage and use the external table to load the data into Snowflake.

C. Copy the data from Google Cloud Storage to Azure Blob storage using external tools and load data from Blob storage to Snowflake.

D. Create a Snowflake Account in the Google Cloud Platform (GCP), ingest data into this account and use data replication to move the data from GCP to Azure.

Suggested Answer: *B*

Community vote distribution

B (100%)

Soumak Most Recent 1 year, 8 months ago

this only looks most viable.

upvoted 2 times

67422cb 1 year, 10 months ago

I wish these questions would be worded better. It sounds like you are supposed to create an external stage in Google Cloud Storage, not in Snowflake referencing Google Cloud Storage. But I assume that the latter is what they meant.

upvoted 1 times

Question #75 *Topic 1*

How can the Snowpipe REST API be used to keep a log of data load history?

A. Call insertReport every 20 minutes, fetching the last 10,000 entries.

B. Call loadHistoryScan every minute for the maximum time range.

C. Call insertReport every 8 minutes for a 10-minute time range.

D. Call loadHistoryScan every 10 minutes for a 15-minute time range.

Suggested Answer: *D*

Community vote distribution

C (100%)

MSIDDIQUI18 Most Recent 1 year, 10 months ago

Correct Answer: C "loadHistoryScan" Snowflake REST API endpoint is rate limited to avoid excessive calls. To help avoid exceeding the rate limit (error code 429), we recommend relying more heavily on insertReport than loadHistoryScan.

upvoted 1 times

akellaanurag 1 year, 11 months ago

Correct answer is C Important This endpoint is rate limited to avoid excessive calls. To help avoid exceeding the rate limit (error code 429), we recommend relying more heavily on insertReport than loadHistoryScan. When calling loadHistoryScan, specify the most narrow time range that includes a set of data loads. For example, reading the last 10 minutes of history every 8 minutes would work well. Trying to read the last 24 hours of history every minute will result in 429 errors indicating a rate limit has been reached. The rate limits are designed to allow each history record to be read a handful of times.

upvoted 2 times

Acaer 2 years, 1 month ago

By calling this command repeatedly, it is possible to see the full history of events on a pipe over time. Note that the command must be called often enough to not miss events. Events are retained for a maximum of 10 minutes. https://docs.snowflake.com/en/user-guide/data-load-snowpipe-rest-apis#endpoint-insertreport

upvoted 3 times

cui_li 2 years, 6 months ago

Answer should be C. "..To help avoid exceeding the rate limit (error code 429), we recommend relying more heavily on insertReport than loadHistoryScan.." https://docs.snowflake.com/en/user-guide/data-load-snowpipe-rest-apis

upvoted 2 times

Atomic_Gecko 2 years, 7 months ago

C is correct - insertReport should be favored over loadHistory as excessive usage of the latter tends to lead to API throttling

upvoted 2 times

Question #76 *Topic 1*

The diagram shows the process flow for Snowpipe auto-ingest with Amazon Simple Notification Service (SNS) with the following steps:

Step 1: Data files are loaded in a stage.

Step 2: An Amazon S3 event notification, published by SNS, informs Snowpipe — by way of Amazon Simple Queue Service (SQS) - that files are ready to load. Snowpipe copies the files into a queue.

Step 3: A Snowflake-provided virtual warehouse loads data from the queued files into the target table based on parameters defined in the specified pipe.

If an AWS Administrator accidentally deletes the SQS subscription to the SNS topic in Step 2, what will happen to the pipe that references the topic to receive event messages from Amazon
S3?

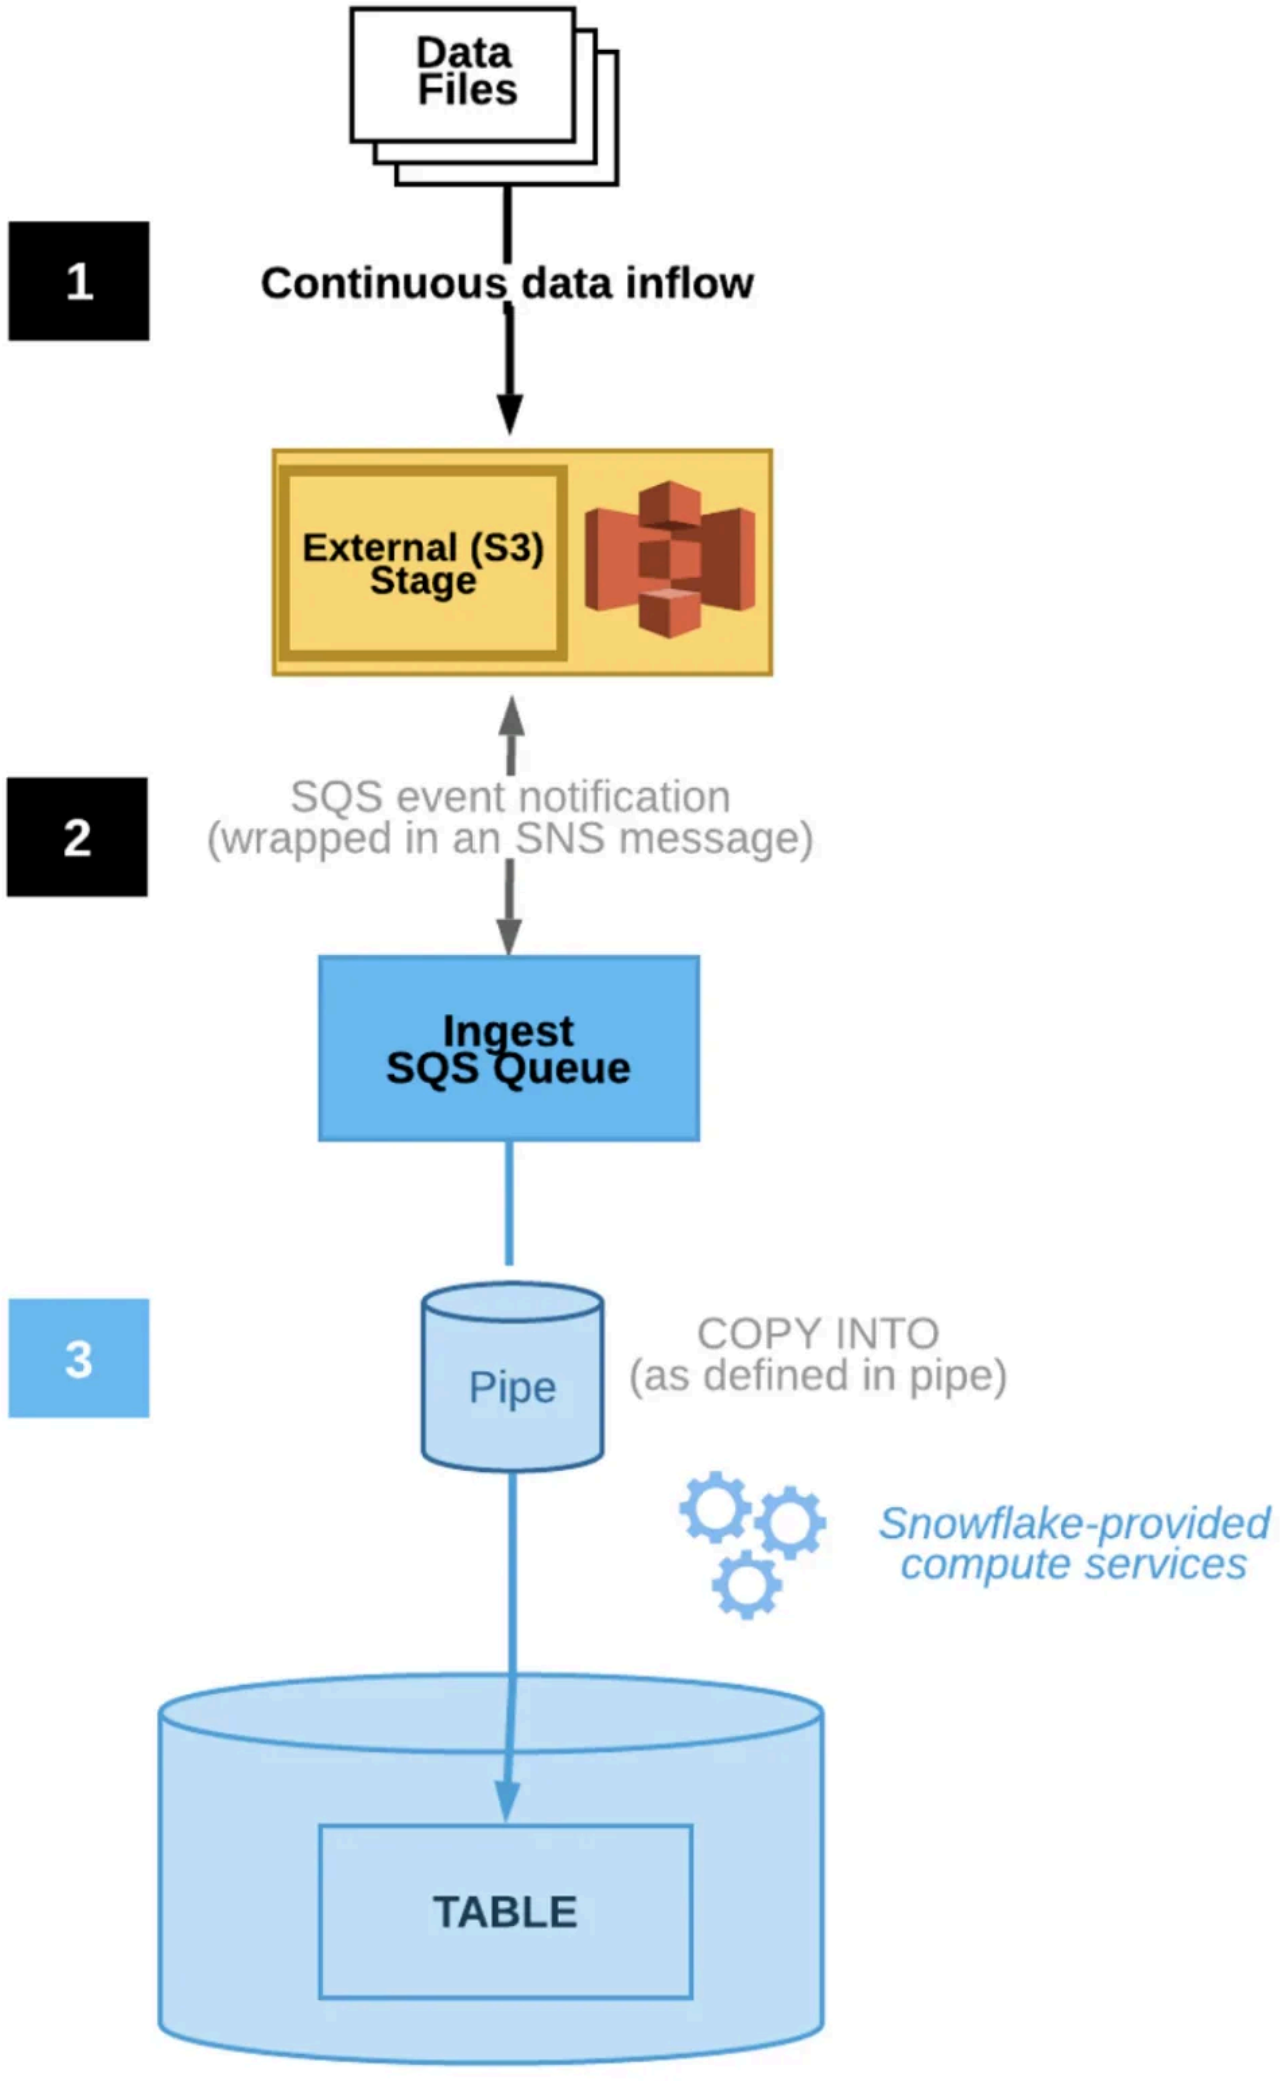
Data Files
1
Continuous data inflow
External (S3) Stage
2
SQS event notification
(wrapped in an SNS message)
Ingest SQS Queue
3
Pipe
COPY INTO
(as defined in pipe)
Snowflake-provided compute services
TABLE

A. The pipe will continue to receive the messages as Snowflake will automatically restore the subscription to the same SNS topic and will recreate the pipe by specifying the same SNS topic name in the pipe definition.

B. The pipe will no longer be able to receive the messages and the user must wait for 24 hours from the time when the SNS topic subscription was deleted. Pipe recreation is not required as the pipe will reuse the same subscription to the existing SNS topic after 24 hours.

C. The pipe will continue to receive the messages as Snowflake will automatically restore the subscription by creating a new SNS topic. Snowflake will then recreate the pipe by specifying the new SNS topic name in the pipe definition.

D. The pipe will no longer be able to receive the messages. To restore the system immediately, the user needs to manually create a new SNS topic with a different name and then recreate the pipe by specifying the new SNS topic name in the pipe definition.

Suggested Answer: *D*

Community vote distribution



Arkitekt Most Recent 1 year, 5 months ago

"To circumvent the 72-hour delay, you can create a SNS topic with a different name. Recreate any pipes that reference the topic using the CREATE OR REPLACE PIPE command, and specify the new topic name." https://docs.snowflake.com/en/user-guide/data-load-snowpipe-ts#snowpipe-stops-loading-files-after-amazon-sns-topic-subscription-is-deleted

upvoted 2 times

ywan1600 1 year, 7 months ago

D is correct https://docs.snowflake.com/en/user-guide/data-load-snowpipe-ts#snowpipe-stops-loading-files-after-amazon-sns-topic-subscription-is-deleted

upvoted 2 times

swe12 2 years ago

D is right https://docs.snowflake.com/en/user-guide/data-load-snowpipe-ts

upvoted 1 times

Question #77 *Topic 1*

An Architect needs to meet a company requirement to ingest files from the company's AWS storage accounts into the company's Snowflake Google Cloud Platform (GCP) account.

How can the ingestion of these files into the company's Snowflake account be initiated? (Choose two.)

A. Configure the client application to call the Snowpipe REST endpoint when new files have arrived in Amazon S3 storage.

B. Configure the client application to call the Snowpipe REST endpoint when new files have arrived in Amazon S3 Glacier storage.

C. Create an AWS Lambda function to call the Snowpipe REST endpoint when new files have arrived in Amazon S3 storage.

D. Configure AWS Simple Notification Service (SNS) to notify Snowpipe when new files have arrived in Amazon S3 storage.

E. Configure the client application to issue a COPY INTO command to Snowflake when new files have arrived in Amazon S3 Glacier storage.

Suggested Answer: *A, C*

Community vote distribution

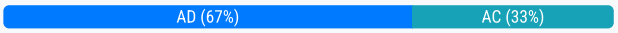

Arkitekt Most Recent 1 year, 5 months ago

I don't think D completely covers the requirements in the question. The way it is stated, the SNS only notifies when new files have arrived. It doesn't initiate the ingestion of the files.

upvoted 1 times

ukpino 1 year, 9 months ago

I think A and D os best. What is the advantage of setting up lambda instead of using SNS, which is recommended for Snowpipe?

upvoted 2 times

MSIDDIQUI18 2 years, 3 months ago

A: This approach involves configuring a client application to detect new files in Amazon S3 and call the Snowpipe REST API to trigger the data loading process into Snowflake. This method provides flexibility and control over the ingestion process. C: Using AWS Lambda to automate the process of detecting new files in Amazon S3 and calling the Snowpipe REST API is an efficient and serverless solution. This approach leverages AWS Lambda's event-driven architecture to trigger Snowpipe ingestion automatically.

upvoted 3 times

Question #78 *Topic 1*

When activating Tri-Secret Secure in a hierarchical encryption model in a Snowflake account, at what level is the customer-managed key used?

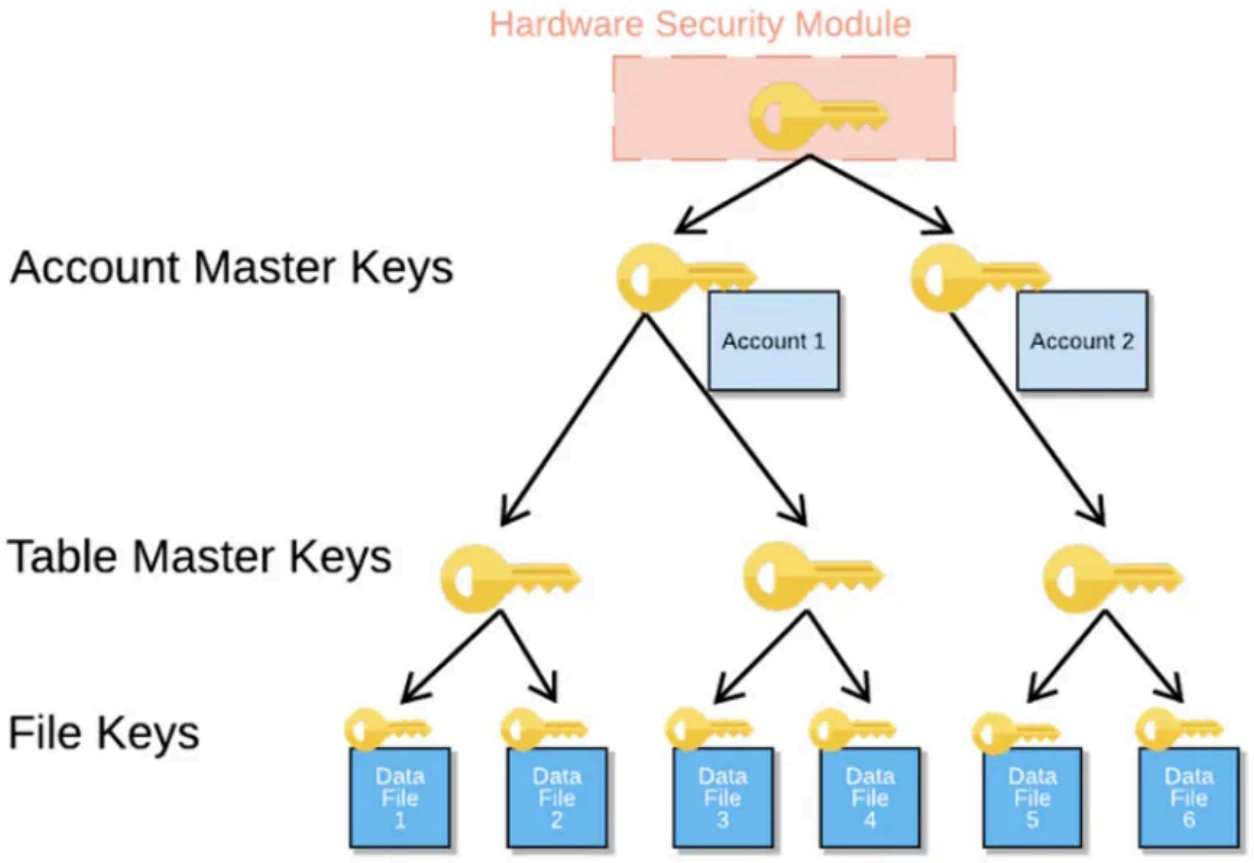

A. At the root level (HSM)

B. At the account level (AMK)

C. At the table level (TMK)

D. At the micro-partition level

Suggested Answer: *B*

Community vote distribution

B (100%)

ywan1600 Most Recent 1 year, 7 months ago

B is correct

upvoted 1 times

isadd_ 2 years, 2 months ago

B, at the account level

upvoted 1 times

Question #79 *Topic 1*

What are characteristics of the use of transactions in Snowflake? (Choose two.)

A. Explicit transactions can contain DDL, DML, and query statements.

B. The AUTOCOMMIT setting can be changed inside a stored procedure.

C. A transaction can be started explicitly by executing a BEGIN WORK statement and end explicitly by executing a COMMIT WORK statement.

D. A transaction can be started explicitly by executing a BEGIN TRANSACTION statement and end explicitly by executing an END TRANSACTION statement.

E. Explicit transactions should contain only DML statements and query statements. All DDL statements implicitly commit active transactions.

Suggested Answer: *C, E*

Community vote distribution

CE (100%)

Arkitekt Most Recent 1 year, 5 months ago

"A. Explicit transactions should contain only DML statements and query statements. DDL statements implicitly commit active transactions B. Caution: Do not change AUTOCOMMIT settings inside a stored procedure. You will get an error message. C. A transaction can be started explicitly by executing a BEGIN statement. Snowflake supports the synonyms BEGIN WORK and BEGIN TRANSACTION. Snowflake recommends using BEGIN TRANSACTION. A transaction can be ended explicitly by executing COMMIT or ROLLBACK. Snowflake supports the synonym COMMIT WORK for COMMIT, and the synonym ROLLBACK WORK for ROLLBACK. D. BEGIN TRANSACTION can be used to start a transaction, but there is no END TRANSACTION statement. To end, COMMIT or ROLLBACK is used. E. (Same note as A) Explicit transactions should contain only DML statements and query statements. DDL statements implicitly commit active transactions https://docs.snowflake.com/en/sql-reference/transactions#what-is-a-transaction"

upvoted 1 times

Soumak 1 year, 8 months ago

validated https://docs.snowflake.com/en/sql-reference/transactions

upvoted 2 times

swe12 2 years ago

CE is right https://docs.snowflake.com/en/sql-reference/transactions

upvoted 2 times

Question #80 *Topic 1*

Which query will identify the specific days and virtual warehouses that would benefit from a multi-cluster warehouse to improve the performance of a particular workload?

```
SELECT TO_DATE(START_TIME) AS DATE,
WAREHOUSE_NAME,
BYTES_SCANNED,
BYTES_SPILLED
FROM "SNOWFLAKE"."ACCOUNT_USAGE"."QUERY_HISTORY"
HAVING BYTES_SPILLED>BYTES_SCANNED;
```

```
SELECT TO_DATE(START_TIME) AS DATE,
WAREHOUSE_NAME,
SUM(AVG_RUNNING) AS SUM_RUNNING,
SUM(AVG_QUEUED_LOAD) AS SUM_QUEUED
FROM "SNOWFLAKE"."ACCOUNT_USAGE"."WAREHOUSE_LOAD_HISTORY"
GROUP BY 1,2
HAVING SUM(AVG_QUEUED_LOAD) > 0;
```

```
SELECT TO_DATE(START_TIME) AS DATE,
WAREHOUSE_NAME,
BYTES_SCANNED,
BYTES_SPILLED
FROM "SNOWFLAKE"."ACCOUNT_USAGE"."WAREHOUSE_LOAD_HISTORY"
HAVING BYTES_SPILLED>BYTES_SCANNED;
```

```
SELECT TO_DATE(START_TIME) AS DATE,
WAREHOUSE_NAME,
BYTES_SPILLED_TO_LOCAL_STORAGE,
SUM(AVG_QUEUED_LOAD) AS SUM_QUEUED
FROM "SNOWFLAKE"."ACCOUNT_USAGE"."QUERY_HISTORY"
HAVING BYTES_SPILLED_TO_LOCAL_STORAGE >0;
```

A.

B.

C.

D.

Suggested Answer: *B*

Community vote distribution

B (100%)

NachoPrendes Highly Voted 2 years, 2 months ago

I think B, because is the only one that uses the HAVING clause correctly (with an aggregate). The rest ones would use WHERE instead of HAVING

upvoted 5 times

MSIDDIQUI18 Most Recent 1 year, 10 months ago

Correct Answer "B" The "Account Usage" view has following columns, "START_TIME, END_TIME, WAREHOUSE_ID, WAREHOUSE_NAME, AVG_RUNNING, AVG_QUEUED_LOAD, AVG_QUEUED_PROVISIONING, AVG_BLOCKED"

upvoted 1 times

cui_li 2 years, 4 months ago

Should D.

upvoted 1 times

Question #81 *Topic 1*

A company is designing high availability and disaster recovery plans and needs to maximize redundancy and minimize recovery time objectives for their critical application processes.

Cost is not a concern as long as the solution is the best available.

The plan so far consists of the following steps:

1. Deployment of Snowflake accounts on two different cloud providers.
2. Selection of cloud provider regions that are geographically far apart.
3. The Snowflake deployment will replicate the databases and account data between both cloud provider accounts.
4. Implementation of Snowflake client redirect.

What is the MOST cost-effective way to provide the HIGHEST uptime and LEAST application disruption if there is a service event?

A. Connect the applications using the - URL. Use the Business Critical Snowflake edition.

B. Connect the applications using the - URL. Use the Virtual Private Snowflake (VPS) edition.

C. Connect the applications using the - URL. Use the Enterprise Snowflake edition.

D. Connect the applications using the - URL. Use the Business Critical Snowflake edition.

Suggested Answer: *A*

Community vote distribution

A (100%)

Soumak Most Recent 1 year, 8 months ago

Both A and D are correct

upvoted 1 times

cui_li 1 year, 10 months ago

Both A and D are identical....both are correct. But can only select one... something isn't right here

upvoted 2 times

Question #82 *Topic 1*

A company's Architect needs to find an efficient way to get data from an external partner, who is also a Snowflake user. The current solution is based on daily JSON extracts that are placed on an FTP server and uploaded to Snowflake manually. The files are changed several times each month, and the ingestion process needs to be adapted to accommodate these changes.

What would be the MOST efficient solution?

A. Ask the partner to create a share and add the company's account.

B. Ask the partner to use the data lake export feature and place the data into cloud storage where Snowflake can natively ingest it (schema-on-read).

C. Keep the current structure but request that the partner stop changing files, instead only appending new files.

D. Ask the partner to set up a Snowflake reader account and use that account to get the data for ingestion.

Suggested Answer: *A*

Community vote distribution



Anblicks Most Recent 1 year, 2 months ago

data sharing is easy.

upvoted 1 times

Soumak 1 year, 8 months ago

already the other user is a snowflake user, so share would be the optimal solution.

upvoted 2 times

AleksM 1 year, 9 months ago

json is prepared from some table(tables) and extracts to FTP. they can create a share with secure view which will contain all data with logic in place

upvoted 2 times

NachoPrendes 2 years, 2 months ago

A. No, users sharing a share that includes a database cannot see the files contained in an stage within that database by default.

upvoted 2 times

Based on the architecture in the image, how can the data from DB1 be copied into TBL2? (Choose two.)

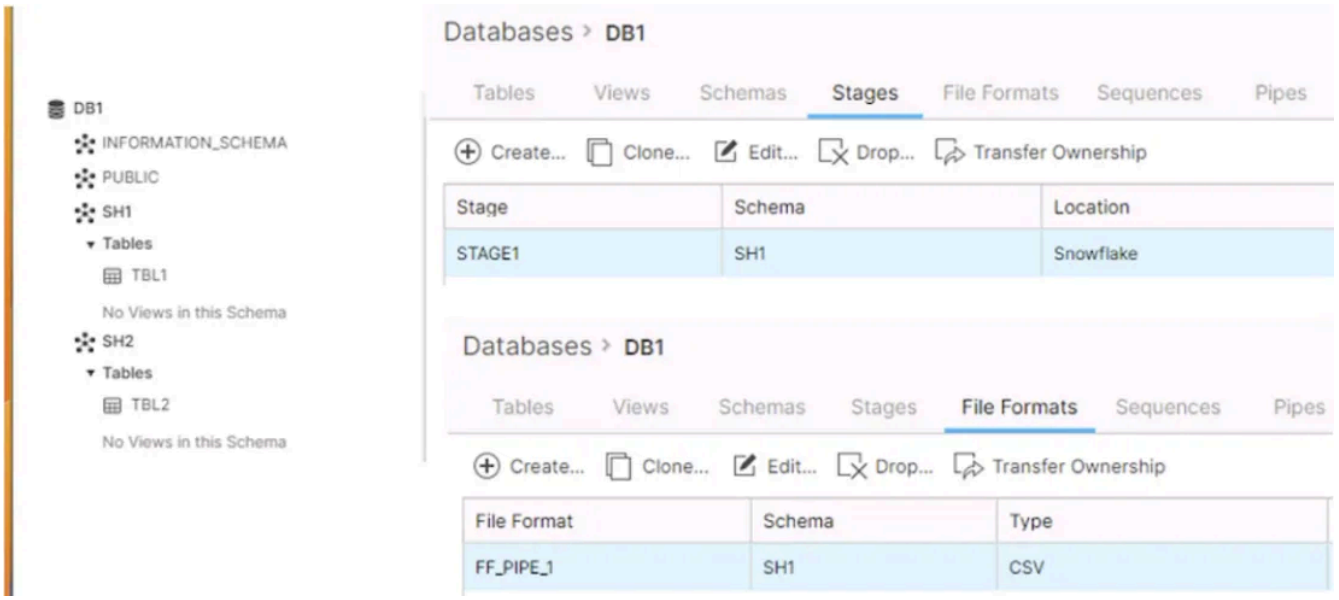


```
use database DB1;
use schema SH1;
copy into DB1.SH2.TBL2
  from @STAGE1
 file_format = (format_name = FF_PIPE_1);

use database DB1;
use schema SH2;
copy into TBL2
  from @STAGE1
 file_format = (format_name = FF_PIPE_1);

use database DB1;
use schema SH2;
copy into DB1.SH2.TBL2
  from @STAGE1
 file_format = (format_name = DB1.SH1.FF_PIPE_1);

use database DB1;
use schema SH2;
copy into DB1.SH2.TBL2
  from @DB1.SH1.STAGE1
 file_format = (format_name = FF_PIPE_1);

use database DB1;
use schema SH2;
copy into TBL2
  from @DB1.SH1.STAGE1
 file_format = (format_name = DB1.SH1.FF_PIPE_1);
```

A.

B.

C.

D.

E.

Suggested Answer: *C, E*

Community vote distribution

AE (100%)

babuka_94 Most Recent 1 year, 8 months ago

A & E are right answers.

upvoted 2 times

Soumak 1 year, 8 months ago

AE should be the correct as the have the complete paths

upvoted 1 times

MSIDDIQUI18 1 year, 10 months ago

A & E are correct answers. C is missing complete path for @STAGE, since schema is using SH2 and stage is in schema DB1.SH1.STAGE1

upvoted 1 times

Pk2287 2 years, 2 months ago

A & E are right answers.

upvoted 2 times

cui_li 2 years, 4 months ago

A & E should be correct.

upvoted 2 times

Atomic_Gecko 2 years, 7 months ago

A & E are the only ones with qualified names that will work

upvoted 3 times

Question #84 *Topic 1*

Why might a Snowflake Architect use a star schema model rather than a 3NF model when designing a data architecture to run in Snowflake? (Choose two.)

A. Snowflake cannot handle the joins implied in a 3NF data model.

B. The Architect wants to remove data duplication from the data stored in Snowflake.

C. The Architect is designing a landing zone to receive raw data into Snowflake.

D. The BI tool needs a data model that allows users to summarize facts across different dimensions, or to drill down from the summaries.

E. The Architect wants to present a simple flattened single view of the data to a particular group of end users.

Suggested Answer: *D, E*

Community vote distribution

DE (100%)

Soumak Most Recent 1 year, 8 months ago

agreed with NachoPrendes

upvoted 1 times

NachoPrendes 2 years, 2 months ago

A. Snowflake handle the joins perfectly B. data duplication is bigger in Star Schema C. The landing zone receives the data "AS IS", without changes

upvoted 1 times

Question #85 *Topic 1*

An Architect is troubleshooting a query with poor performance using the QUERY_HISTORY function. The Architect observes that the COMPILATION_TIME is greater than the EXECUTION_TIME.

What is the reason for this?

A. The query is processing a very large dataset.

B. The query has overly complex logic.

C. The query is queued for execution.

D. The query is reading from remote storage.

Suggested Answer: *B*

Community vote distribution

B (100%)

Soumak Most Recent 1 year, 8 months ago

unless it is a complex query, compilation time would not increase

upvoted 2 times

Question #86 *Topic 1*

A Snowflake Architect is designing a multiple-account design strategy.

This strategy will be MOST cost-effective with which scenarios? (Choose two.)

A. The company wants to clone a production database that resides on AWS to a development database that resides on Azure.

B. The company needs to share data between two databases, where one must support Payment Card Industry Data Security Standard (PCI DSS) compliance but the other one does not.

C. The company needs to support different role-based access control features for the development, test, and production environments.

D. The company security policy mandates the use of different Active Directory instances for the development, test, and production environments.

E. The company must use a specific network policy for certain users to allow and block given IP addresses.

Suggested Answer: *B, D*

Community vote distribution

BD (100%)

ywan1600 Most Recent 1 year, 7 months ago

I'd go with B and D. Option C is doable within a single tenant, and it might even save you money in certain situations, like zero-copy cloning.

upvoted 1 times

MSIDDIQUI18 2 years, 3 months ago

C: Using separate Snowflake accounts for development, test, and production environments allows for tailored RBAC configurations without intermingling roles and permissions. This separation can simplify access management and enhance security, leading to cost-effective administration and compliance with organizational policies. D: If the security policy requires different AD instances for various environments, maintaining separate Snowflake accounts is essential. This segregation ensures that each environment adheres to its specific security requirements, reducing complexity and potential security risks, thus optimizing costs related to security management and compliance.

upvoted 2 times

Question #87 *Topic 1*

The following table exists in the production database:

A regulatory requirement states that the company must mask the username for events that are older than six months based on the current date when the data is queried.

How can the requirement be met without duplicating the event data and making sure it is applied when creating views using the table or cloning the table?

Name: `user_events`

Column	Description
`user_hash_key`	hashed key with no meaning
`username`	user name
`event_type`	the type of event that occurred
`event_timestamp`	the date and time the event occurred

A. Use a masking policy on the username column using an entitlement table with valid dates.

B. Use a row level policy on the user_events table using an entitlement table with valid dates.

C. Use a masking policy on the username column with event_timestamp as a conditional column.

D. Use a secure view on the user_events table using a case statement on the username column.

Suggested Answer: *C*

Community vote distribution

C (75%) A (25%)

BobFar Most Recent 11 months, 2 weeks ago

C is the best answer. A and B involve extra complexity and/or duplication of data. D is technically feasible, but it is less robust and enforceable compared to a masking policy, users could still bypass the secure view and query the underlying table directly.

upvoted 1 times

ywan1600 1 year, 7 months ago

C is correct, as ukpino explained in his second comment.

upvoted 2 times

ukpino 1 year, 9 months ago

Selected Answer: A

C does not seem right, because event_timestamp does not matter, current date being queried does. A seems more accurate.

upvoted 1 times

ukpino 1 year, 9 months ago

C does not seem right, because event_timestamp does not matter, current date being queried does. A seems more accurate.Actually changing to C, event-timestamp would be the column to be compared against the current datetime, C makes sense

upvoted 1 times

ukpino 1 year, 9 months ago

Actually changing to C, event-timestamp would be the column to be compared against the current datetime, C makes sense

upvoted 1 times

Question #88 *Topic 1*

Which data models can be used when modeling tables in a Snowflake environment? (Choose three.)

A. Graph model

B. Dimensional/Kimball

C. Data lake

D. Inmon/3NF

E. Bayesian hierarchical model

F. Data vault

Suggested Answer: *B, D, F*

Community vote distribution

BDF (100%)

ywan1600 Most Recent 1 year, 7 months ago

BDF. The others are not data modelling methods..

upvoted 1 times

MSIDDIQUI18 2 years, 3 months ago

B: Dimensional Modeling (Kimball): This approach organizes data into fact and dimension tables. It's widely used in data warehousing to optimize for query performance and ease of understanding for end-users. Snowflake supports this model well due to its columnar storage and performance optimization features. D: Inmon/3NF: This model, proposed by Bill Inmon, focuses on a normalized structure (Third Normal Form) which is more flexible and reduces redundancy. Snowflake's capabilities support the normalized model, ensuring data integrity and reducing storage costs. F: Data Vault: This modeling approach is designed to provide long-term historical storage of data from multiple operational systems. It supports agile and scalable data integration and changes. Snowflake's support for semi-structured data and its efficient storage makes it a good fit for the Data Vault methodology.

upvoted 2 times

Question #89 *Topic 1*

A Snowflake Architect is setting up database replication to support a disaster recovery plan. The primary database has external tables.

How should the database be replicated?

A. Create a clone of the primary database then replicate the database.

B. Move the external tables to a database that is not replicated, then replicate the primary database.

C. Replicate the database ensuring the replicated database is in the same region as the external tables.

D. Share the primary database with an account in the same region that the database will be replicated to.

Suggested Answer: *B*

Community vote distribution

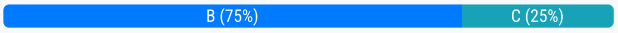

785de4c Most Recent 1 year, 1 month ago

https://docs.snowflake.com/en/release-notes/bcr-bundles/2024_02/bcr-1528

upvoted 1 times

ywan1600 1 year, 7 months ago

B. You can't replicate external tables.

upvoted 2 times

MultiCloudIronMan 1 year, 7 months ago

C. Replicate the database ensuring the replicated database is in the same region as the external tables. This ensures that the external tables remain accessible and performant since they are in the same region as the primary database.

upvoted 1 times

isadd_ 2 years, 2 months ago

I think it's B, one has to move the external table before replication because otherwise the replication process could rise an error

upvoted 2 times

Question #90 *Topic 1*

An Architect is integrating an application that needs to read and write data to Snowflake without installing any additional software on the application server.

How can this requirement be met?

A. Use SnowSQL.

B. Use the Snowpipe REST API.

C. Use the Snowflake SQL REST API.

D. Use the Snowflake ODBC driver.

Suggested Answer: *C*

Community vote distribution

C (100%)

Soumak Most Recent 1 year, 8 months ago

validated

upvoted 2 times

MSIDDIQUI18 1 year, 10 months ago

Correct Answer: C Snowflake SQL REST API can be used to develop applications and integrations that can READ and WRITE data to Snowflake without installing any additional software on the application server.

upvoted 1 times

MSIDDIQUI18 2 years, 3 months ago

C: Snowflake SQL REST API: This option allows you to execute SQL statements on Snowflake using RESTful HTTP requests. It does not require any additional software to be installed on the application server. You can use standard HTTP methods to interact with Snowflake, making it an ideal choice for applications that need to read and write data to Snowflake without additional installations.

upvoted 1 times

Question #91 *Topic 1*

What transformations are supported in the below SQL statement? (Choose three.)

CREATE PIPE ... AS COPY ... FROM (...)

A. Data can be filtered by an optional WHERE clause.

B. Columns can be reordered.

C. Columns can be omitted.

D. Type casts are supported.

E. Incoming data can be joined with other tables.

F. The ON_ERROR - ABORT_STATEMENT command can be used.

Suggested Answer: *A, B, C*

Community vote distribution

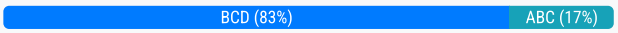

akellaanurag Most Recent 1 year, 11 months ago

correct answer is B,C,D Using a query as the source for the COPY statement for column reordering, column omission, 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.

upvoted 1 times

Acaer 2 years ago

Using a query as the source for the COPY statement for column reordering, column omission, and casts (i.e. transforming data during a load) is supported. https://docs.snowflake.com/en/sql-reference/sql/create-pipe#usage-notes

upvoted 4 times

6c79d6f 2 years ago

Using a query as the source for the COPY statement for column reordering, column omission, and casts (i.e. transforming data during a load) is supported.

upvoted 1 times

imdashi 2 years, 2 months ago

B,C,D, where clause not supported in COPY

upvoted 1 times

NachoPrendes 2 years, 2 months ago

Agree with MSIDDIQUI18

upvoted 1 times

MSIDDIQUI18 2 years, 3 months ago

A: You can specify a WHERE clause to filter data before it is copied into Snowflake. This allows you to select specific rows based on certain conditions. B: When you define the structure of the target table in Snowflake using CREATE PIPE, you can specify the order of columns. This means you can control the order in which incoming data columns are mapped to the target table. C: Similar to reordering columns, you can specify only the columns from the incoming data that you want to copy into the target table. Columns that are not specified in the CREATE PIPE statement will be ignored during the data copy process.

upvoted 2 times

Atomic_Gecko 2 years, 7 months ago

B, C, D. You can't use where clauses in a copy into statement

upvoted 2 times