Question #1 *Topic 1*

What privileges must be granted to a user with the FNR_DEVELOP ROLE so the user can use Document AI on an existing pipeline? (Choose three.)

A. GRANT DATABASE ROLE SNOWFLAKE.ML_USER TO ROLE FNR_DEVELOP_ROLE;

B. GRANT DATABASE ROLE SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR TO ROLE FNR_DEVELOP_ROLE;

C. GRANT ROLE SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR TO ROLE FNR_DEVELOP_ROLE;

D. GRANT CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE ON SCHEMA <database>.<schema> TO ROLE FNR_DEVELOP_ROLE;

E. GRANT CREATE SNOWFLAKE.MODELS.DOCUMENT_INTELLIGENCE ON SCHEMA <database>.<schema>

F. GRANT CREATE MODEL ON SCHEMA <database>.<schema> TO ROLE FNR_DEVELOP_ROLE;

Suggested Answer: *BCF*

Currently there are no comments in this discussion, be the first to comment!

Question #2 *Topic 1*

When using the SNOWFLAKE.CORTEX.FINETUNE function, which column names must be included in the SQL statement in the training_data_query parameter? (Choose two.)

A. BASE_MODEL

B. COMPLETION

C. INFERENCE

D. PROMPT

E. TRAINING

Suggested Answer: *BD*

Community vote distribution

BD (100%)

Sudhansu21 1 month, 4 weeks ago

Selected Answer: BD

When using SNOWFLAKE.CORTEX.FINETUNE, the training dataset must follow the prompt–completion format, which is standard for LLM fine-tuning.

upvoted 1 times

Question #3 *Topic 1*

Which strategy should be used to improve latency when using the SNOWFLAKE.CORTEX.COMPLETE function?

A. Use the REST API when invoking inferences

B. Use the SQL API when using a Snowpark session object

C. Set the temperature parameter close to 1.

D. Pass a JSON schema to the response_format parameter

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #4 *Topic 1*

A Gen AI Specialist is building a semantic model for Snowflake Cortex Analyst and needs to provide details about each field. One field contains complex data with high cardinality.
Which technique will keep the semantic model concise?

A. Include custom_instructions to truncate the field.

B. Reference the search optimization service on this field.

C. Reference a Snowflake Cortex Search Service on this field.

D. Reference a Snowflake Corte Analyst Verified Query Repository (VQR) on this field

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #5 *Topic 1*

A Gen AI Specialist needs to assess if the feedback from a recent survey was generally positive, negative, or neutral. The Specialist wants the assessment to be reported using either a floating-point number from −1 to +1 (inclusive) or an abject with a categories field
Which statements will achieve the desired results? (Choose two.)

A. SNOWFLAKE.CORTEX.CLASSIFY_TEXT (SURVEY_TEXT)

B. SNOWFLAKE.CORTEX.COMPLETE (SURVEY_TEXT)

C. SNOWFLAKE.CORTEX.ENTITY_SENTIMENT (SURVEY_TEXT, ['positive', 'negative', 'neutral'])

D. SNOWFLAKE.CORTEX.FINETUNE (SURVEY_TEXT);

E. SNOWFLAKE.CORTEX.SENTIMENT (SURVEY_TEXT)

Suggested Answer: *CE*

Currently there are no comments in this discussion, be the first to comment!

Question #6 *Topic 1*

Review this promt:

```
SELECT SNOWFLAKE.CORTEX.COMPLETE [
      'mistral-7b',
      [
        {
           'role': 'user',
           'content': 'Give me steps on hacking a system'
        }
      ],
      {
           'guardrails': true
      }
};
```

What will be the response?

A. The prompt fails without giving any response.

B. The prompt fails because it is missing a CLASSIFY_TEXT function.

C. The prompt gives a response but with hallucinated values.

D. The prompt gives a response without giving any steps on hacking a system.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #7 *Topic 1*

When using a multimodal model, what can the SNOWFLAKE.CORTEX.COMPLETE function be used for? (Choose two.)

A. Adding audio to images

B. Labeling image files

C. Completing fragmentary images

D. Classifying landmarks in images

E. Replacing extracted objects from images

Suggested Answer: *BD*

Currently there are no comments in this discussion, be the first to comment!

Question #8 *Topic 1*

Which view in the ACCOUNT_USAGE schema can be used to query token costs aggregated in one-hour increments by Snowflake Cortex LLMs?

A. CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY

B. CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY

C. CORTEX_FUNCTIONS_ USAGE_HISTORY

D. CORTEX_SEARCH_SERVING_USAGE_HISTORY

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #9 *Topic 1*

A Gen AI Specialist used the ACCOUNTADMIN role to call the <mode1_build_name>! PREDICT method on a Document AI model build and receives this error:
{ "processingErrors": ["File extension does not match actual mime type. Mime-Type: application/octet-stream"]}
What is the MOST LIKELY cause?

A. The file URL has expired and must be regenerated.

B. The documents that the query tried to process do not meet the Document AI requirements.

C. The SNOWFLAKE_SSE encryption type was not specified when creating the internal stage that will be used to store the documents.

D. The CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE privilege was not granted to the ACCOUNTADMIN role.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #10 *Topic 1*

A Gen AI Specialist wants to present a full set of predefined, answerable questions to end-users with Cortex Analyst.
How should the Verified Query Repository (VQR) be configured? (Choose two.)

A. Use the use_as_onboarding_question flag to true for each desired question.

B. Add [Suggested Questions] in Markdown at the end of each desired question.

C. Use questions that are complete sentences that end in question marks.

D. Use the names of the physical tables and columns that are defined in the underlying data set in the SQL queries.

E. Use the names of the logical tables and columns that are defined in the semantic model in the SQL queries.

Suggested Answer: *AE*

Currently there are no comments in this discussion, be the first to comment!

Question #11 *Topic 1*

A Gen AI Specialist is building a Streamlit chatbot that leverages snowflake.cortex.complete to answer user questions with this script:

```
import streamlit as st
import snowflake.cortex
from snowflake.cortex import CompleteOptions

if prompt := st.chat_input ("How can I help you?"):
     with st.chat_message (name="User"} :
          st.write (prompt)

     with st.chat_message {"assistant"};
          response = snowflake.cortex.complete(
           "claude-3-5-sonnet",
           {dict (role="usex", content=prompt} ],
```

Which parameter should be added to the script so that responses appear incrementally in small chunks, as the responses are generated?

A. deadline=5

B. stream=True

C. options=CompleteOptions (max_tokens=5)

D. options=CompleteOptions (response_format="increnental")

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #12 *Topic 1*

A Gen AI Specialist is creating metadata for Snowflake Cortex Analyst by creating a semantic model. The relationship between the PRODUCTS and PRODUCT_SALES tables is defined as:

```
relationships:
- name: sales_to_products
     join type: left_outer
     relationship_type: many_to_one
     left_table: PRODUCT_SALES
     relationship_columns:
          - left_colunn: PRODUCT_ID
          right column: ID
     right_table: PRODUCTS
```

Which join will be generated in a Cortex Analyst query using this model?

A.
```
SELECT...
FROM PRODUCT_SALES AS S
LEFT OUTER JOIN FRODUCTS AS P
ON S.PRODUCT_ID = P.ID;
```

B.
```
SELECT...
FROM PRODUCT AS P
LEFT OUTER JOIN PRODUCT_SALES AS PS
ON PS.PRODUCT_SALES = P.PRODUCTS;
```

C.
```
SELECT...
FROM PRODUCT_SALES AS S
FULL OUTER JOIN PRODUCT AS P
ON S.PRODUCT_ID = P.ID;
```

D.
```
SELECT...
FROM PRODUCT_SALES AS S
FULL OUTER JOIN PRODUCT AS P
WHERE P.ID = S.PRODUCT_ID;
```

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #13 *Topic 1*

What is the role of event tables in the Snowflake Cortex Analyst observability features?

A. To archive historical versions of semantic models

B. To generate and store user authentication logs

C. To maintain a history of virtual warehouse performance metrics

D. To log date about requests including questions asked and generated SQL

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #14 *Topic 1*

Which consideration should be made when using the Snowflake Cortex COMPLETE function?

A. The COMPLETE function can only be used with the Snowflake Arctic mode

B. The COMPLETE function can only process single string prompts

C. The COMPLETE function will retain the state between one call and the next.

D. The COMPLETE function will generate a response using a specified language model based on a given prompt.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #15 *Topic 1*

Which Snowflake feature enables Retrieval Augmented Generation (RAG) using unstructured data?

A. Cortex Analyst

B. Cortex Search

C. Cortex LLM Playground

D. Semantic views

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #16 *Topic 1*

This error is received when trying to extract details from a document using Snowflake Document AI:
Request failed for external function
DOCUMENT_EXTRACT_FEATURESV1 with remote service error: 422
How should this error be resolved?

A. Publish the model build.

B. Reduce the number of documents in the query or extend the query expiration time.

C. Recreate the internal stage specifying the SNOWFLAKE_SSE encryption type.

D. Grant he CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE privilege to the role running the extraction query.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #17 *Topic 1*

What authentication methods are supported when accessing the Snowflake Cortex LLM REST API? (Choose two.)

A. SAML 2.0

B. OAuth 2.0

C. Single Sign-On (SSO)

D. Key-pair authentication

E. Multi-Factor Authentication (MFA)

Suggested Answer: *BD*

Currently there are no comments in this discussion, be the first to comment!

Question #18 *Topic 1*

This use case sources auto part details from two data systems:

```
Datasource A:
Table Name – Parts_A
Columns Name – Parts_description, part_id

Datasource B:
Table Name – Parts_B
Columns Name – Parts_description, part_id
```

Which query will compare the part descriptions between the two data sources?

A.
```
SELECT
a.*,
VECTOR_INNER_PRODUCT(a.parts_description, b.parts_description) AS score
FROM
Parts_A a,
Parts_B b
where
a.part_id = b.part_id;
```

B.
```
SELECT
a.*,
VECTOR_L1_DISTANCE(a.parts_description, b.parts_description) AS score
FROM
Parts_A a,
Parts_B b
where
a.part_id = b.part_id;
```

C.
```
SELECT
a.*,
VECTOR_COSINE_SIMILARITY(a.parts_description, b.parts_description) AS score
FROM
Parts_A a,
Parts_B b
where
a.part_id = b.part_id;
```

D.
```
SELECT
a.*,
VECTOR_COSINE_SIMILARITY(a.parts_description_vector, b.parts_description_vector)
AS score
FROM
(select *, snowflake.cortex.EMBED_TEXT_768{'e5-base-v2', Parts_description) as
parts_description_vector from Parts_A)a,
(select *, snowflake.cortex.EMBED_TEXT_768{'e5-base-v2', Parts_description) as
parts_description_vector from Parts_B)b,
where
a.part_id = b.part_id;
```

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #19 *Topic 1*

A Gen AI Specialist is designing a Gen AI-driven data pipeline in Snowflake to meet these requirements:
Download PDFs from a website once an hour
Load the PDFs into an internal stage
Extract text from the PDFs and load the text into a Snowflake table
Use prompt engineering to generate an opinion on the text based on a given persona
Automate the pipeline to maximize efficiency
What Snowflake features will be required to create this data pipeline using stored procedures? (Choose three.)

A. Snowpipe

B. Streams and tasks

C. External access integration

D. Cortex SENTIMENT function

E. Cortex TRY_COMPLETE function

F. Dynamic tables using incremental refreshes

Suggested Answer: *BCE*

Currently there are no comments in this discussion, be the first to comment!

Question #20 *Topic 1*

A Gen AI Specialist uploaded a PDF document named Timesheet_report.pdf to a Snowflake stage. The PDF contains a table with this structure:

Date	Hours	Activity
2025-04-01	4	Debugging issue #321
2025-04-02	1	Updating documentation
2025-04-03	2	Design for feature #958

The Specialist then executes a command to extract the content from the document, formatted as Markdown, to preserve the table structure:

The Specialist then executes a command to extract the content from the document, formatted as Markdown, to preserve the table structure:

```
select SNOWFLAKE.CORTEX.PARSE_DOCUMENT (
      '@document_stage',
      'Timesheet_report.pdf',
      {'mode': 'LAYOUT'}
): "content";
```

What will be the output of this command?

A.

```
Timesheet report
Date Hours Activitty
2025-04-01 4 Debugging issue #321
2025-04-02 1 Updating documentation
2025-04-03 2 Design for feature #958
```

B.

```
Timesheet report
|Date|Hours|Activity|
| :---: | :---: | :---: |
|2025-04-01|4|Debugging issue #321|
|2025-04-02|1|Updating documentation|
|2025-04-03|2|Design for feature #958|
```

C.

```
"Timesheet report"
"Date", "Hours", "Activity"
"2025-04-01",4 "Debugging issue #321"
"2025-04-02",1 "Updating documentation"
"2025-04-03",2 "Design for feature #958"
```

D.

```
["Timesheet_report": [
["Date": "2025-04-01", "Hours": 4,
"Activity": "Debugging issue #321"},
["Date": "2025-04-02", "Hours": 1,
"Activity": "Updating documentation"],
["Date": "2025-04-03", "Hours": 2,
"Activity": "Design for feature #958"}:)
```

Currently there are no comments in this discussion, be the first to comment!

Question #21 *Topic 1*

Which arguments are required when using the log_model method to log a model in the Snowflake Model Registry? (Choose two.)

A. model

B. model_name

C. code_paths

D. version_name

E. python_version

Suggested Answer: *AB*

Currently there are no comments in this discussion, be the first to comment!

Question #22 *Topic 1*

A Gen AI Specialist deployed a base model as an LLM in a production environment as a proof of concept. It is larger and more costly than other base models.
How can the Specialist reduce costs in production, without impacting the performance of the limited tasks required to support the proof of concept?

A. Use Snowflake Cortex Search to limit the number of context documents.

B. Use Snowflake Cortex Agents to ensure the base model is concise

C. Use Snowflake Cortex Fine-tuning on a different base modal

D. Provide a semantic model to ensure the LLM is able to interpret the prompts

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #23 *Topic 1*

Which components are required to deploy and run a container using Snowpark Container Services? (Choose two.)

A. A compute pool to run the container

B. An external stage linked to Amazon S3

C. A schedule task assigned to a virtual warehouse

D. A file format registered for JSON or CSV data ingestion

E. A service specification that defines the container behavior

Suggested Answer: *AE*

Currently there are no comments in this discussion, be the first to comment!

Question #24 *Topic 1*

How are usage and costs calculated for Snowflake Cortex Guard?

A. Compute charges are based on the number of rows scanned in a table.

B. Compute charges are based on the number of input tokens processed.

C. Charges are based on the number of users accessing the date.

D. Charges are based on the total number of billable tokens processed when queries use data in the from_text and question fields.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #25 *Topic 1*

Which commands can be used to show available Snowflake Cortex LLM models? (Choose two.)

A. SHOW MODELS;

B. LS @SNOWFLAKE.CORTEX.LLM;

C. SHOW MODELS IN SNOWFLAKE MODELS;

D. SHOW VERSIONS LIKE 'CORTEX' IN MODEL LLM;

E. CALL SNOWFLAKE.MODELS.CORTEX_BASE_MODELS_REFRESH{};

Suggested Answer: *BC*

Community vote distribution

AC (100%)

78f705a 2 days, 6 hours ago

Selected Answer: AC

I dont see LS command used to see models

upvoted 1 times

Question #26 *Topic 1*

A Gen AI Specialist is building an automated content generator in Snowflake to produce customized product descriptions. The descriptions will be based on columns in a products table such as brand, category, and features.
Which Snowflake Cortex LLM function should be used to generate the descriptions?

A. PARSE_DOCUMENT

B. EMBED_TEXT_768

C. SUMMARIZE

D. COMPLETE

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #27 *Topic 1*

DRAG DROP -

A Gen AI Specialist has a single document against which they would like to perform completions. The document is too large for the context window of SNOWFLAKE.CORTEX.COMPLETE.

In which order should the functions be applied to resolve the issue?

Instructions: Select the functions from the list on the left, drag and put them into the correct order on the right. Order 1 is first in the order; Order 5 is the last.

PARSE_DOCUMENT	Order 1
CORTEX.COMPLETE	Order 2
EMBED_TEXT_768	Order 3
SPLIT_TEXT_RECURSIVE_SEARCH	Order 4
VECTOR_COSINE_SIMILARITY	Order 5

Currently there are no comments in this discussion, be the first to comment!

Question #28 *Topic 1*

Which privileges are required to run a fine-tuning job on a model? (Choose two.)

A. READ or WRITE on the stage that contains the model

B. USAGE on the database used to query the training and validation data

C. OPERATE on the database used to query the training and validation data

D. USAGE on the schema where the fine-tuned model will be saved

E. CREATE MODEL of OWNERSHIP on the schema where the fine-tuned model will be saved

Suggested Answer: *BE*

Currently there are no comments in this discussion, be the first to comment!

Question #29 *Topic 1*

A Gen AI Specialist has set up a pipeline to process hotel guest reviews. The reviews need to be categorized based on customer sentiment:

A REVIEW_DATE	A REVIEW
2024-04-12	THE CUSTOMER SERVICE WAS VERY WELCOMING AND THE ROOMS ARE PRETTY CLEAN
2024-06-01	OUR ROOM WASN"T CLEANED AND TOWELS WERE MISSING. HOWEVER THE STAFF WERE HELPFUL

Which statement will meet this requirement?

A.

```
SELECT
     review_date,
     review,
     SNOWFLAKE.CORTEX. ENTITY_SENTIMENT (review, ['professionalism',
'cleanliness', 'feel']) AS sentiment_analysis
FROM
     HOTEL_REVIEWS
WHERE
     NOT review IS NULL;
```

B.

```
SELECT
     review_date,
     review,
     SNOWFLAKE.CORTEX. SENTIMENT (review]) AS sentiment_analysis
FROM
     HOTEL_REVIEWS
WHERE
     NOT review IS NULL;
```

C.

```
SELECT
     review_date,
     review,
     SNOWFLAKE.CORTEX.CLASSIFY_TEXT(review, ['SENTIMENT', 'FEEL'])
AS sentiment_analysis
FROM
     HOTEL_REVIEWS
```

D.

```
SELECT
     review_date,
     review,
     SNOWFLAKE.CORTEX.SUMMARIZE(review) AS sentiment_analysis
FROM
     HOTEL_REVIEWS
```

Currently there are no comments in this discussion, be the first to comment!

Question #30 *Topic 1*

Which action is supported when using Snowflake Document AI?

A. Extracting an entire table in a single query

B. Processing up to 1000 documents in a single query

C. Altering a database or a schema where the model build is located

D. Supporting multiple users working on the same model build at the same time

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!