

Snowflake GES-C02

Snowflake Certified SnowPro Specialty - Gen AI

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

- 1. Up to Date products, reliable and verified.**
- 2. Questions and Answers in PDF Format.**



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/ges-c02>

Latest Version: 6.1

Question: 1

A media company's audience-analytics role runs SQL against a semantic view and its underlying tables without error. When an application using the same role sends a question to the Cortex Analyst REST API naming that semantic view, the request is rejected for insufficient privileges. SNOWFLAKE.CORTEX_USER has been revoked from PUBLIC, and the role holds no database roles in the SNOWFLAKE database.

Which grant resolves this?

Response:

- A. The SNOWFLAKE.CORTEX_ANALYST_USER database role, granted to that role
- B. The SNOWFLAKE.CORTEX_AGENT_USER database role, which covers the agents that orchestrate Analyst
- C. OWNERSHIP of the semantic view itself
- D. CREATE SEMANTIC VIEW on the schema, so the role can rebuild the view under its own name

Answer: A

Explanation:

Requests to Cortex Analyst require a role holding SNOWFLAKE.CORTEX_ANALYST_USER or SNOWFLAKE.CORTEX_USER. Access to the data alone does not admit the role to the service.

Question: 2

An environmental consultancy receives technical submissions of 300 to 450 pages. A cover sheet on the first three pages holds the site identifier, permit class and submission date, which must be extracted as named fields. Calling AI_EXTRACT with the full files fails.

Which approach extracts the fields while staying within the functions' documented limits?

Response:

- A. Keep calling AI_EXTRACT on the full files, but cut the request down to the three cover-sheet questions so that each call stays within the limit
- B. Add config => {'scale_factor': 2.0} so that AI_EXTRACT reads the long files more reliably
- C. Run AI_CLASSIFY on the full files to label the permit class
- D. Parse pages one to three with AI_PARSE_DOCUMENT and a 'page_filter', then pass that text to AI_EXTRACT's text argument

Answer: D

Explanation:

AI_EXTRACT takes documents of up to 125 pages, while AI_PARSE_DOCUMENT takes up to 2,000 and can process only chosen pages, whose text then goes through the text argument.

Question: 3

Which of the following are design principles behind Snowflake's approach to generative AI?
(Choose three.)

Response:

- A. Allow unrestricted egress to any external model endpoint teams prefer
- B. Keep the data in Snowflake, so inference happens inside the governance boundary
- C. Govern AI use with the same role-based access control that governs the data
- D. Scale AI compute elastically with demand

Answer: B,C,D

Explanation:

This is the organizing principle the rest follow from: if the data never leaves, the controls that protect it never stop applying.

Access to the AI functions and to the models they may invoke is granted through roles, so an organization manages AI entitlements with the mechanism it already uses rather than a parallel one. Separation of storage and compute means an AI workload can be given more compute when it needs it and none when it does not, without touching the data.

Question: 4

A business analyst wants to ask questions in plain English and get answers computed from the company's Snowflake tables, without writing SQL.

Which Snowflake capability is designed for this?

Response:

- A. Cortex Analyst
- B. COMPLETE
- C. Cortex Search
- D. CLASSIFY_TEXT

Answer: A

Explanation:

Cortex Analyst answers natural-language questions against a semantic model the team authors, which names the tables, measures, dimensions and synonyms it may use. That semantic model is what makes its answers reliable rather than plausible, and authoring it well is most of the work.

Question: 5

A rail operator's track inspectors complete paper forms that combine printed labels, tick boxes and handwritten remarks, and the forms are scanned to PDF. The operations lead assumes staff must first retype each form into a digital template before any Snowflake document function can use it. Which statement corrects this assumption?

Response:

- A. Tick boxes are read from scans, but handwriting is not
- B. Scanned files are accepted once the stage uses server-side encryption, which enables OCR on them
- C. AI_EXTRACT reads checkboxes and handwriting directly from scanned files
- D. Handwritten remarks are read only after AI_TRANSLATE has converted them into typed text first

Answer: C

Explanation:

AI_EXTRACT understands text, tables, checkboxes, handwriting and other visual elements, and it is optimized for scanned as well as digital-born documents.

Question: 6

A wealth management firm's advisers use a Cortex-based assistant over internal research. The firm has licensed a financial news publisher's archive, which the publisher offers on Snowflake Marketplace as a Cortex Knowledge Extension. The firm wants answers grounded in that archive with attribution back to the source articles, without copying the corpus into its own tables or building an ingestion pipeline. The publisher caps how much of its content can be returned each day.

Which approach fits these requirements?

Response:

- A. Use an external access integration so a UDF copies the publisher's articles nightly into a table for search
- B. Load the archive into a table, embed it with AI_EMBED and rank passages with VECTOR_COSINE_SIMILARITY
- C. Train a Cortex Fine-tuning adaptor on the archive so that the answers reflect its articles
- D. Get the Marketplace listing and have the assistant retrieve from the shared Cortex Search service, citing each result's source URL

Answer: D

Explanation:

A Cortex Knowledge Extension is a Cortex Search service shared through a listing, built for RAG over licensed content. The consumer queries it without holding the corpus, the provider's return limit applies, and a source URL column supports attribution.

Question: 7

Which of the following are workloads that Snowflake's generative AI capabilities are well suited to? (Choose three.)

Response:

- A. Sharing a curated dataset with a partner account, with no AI processing involved
- B. Extracting defined values such as totals and dates from PDF documents held on a stage
- C. Answering questions over enterprise data in a chat interface
- D. Assigning categories to inbound customer support tickets

Answer: B,C,D

Explanation:

Pulling named fields out of variable document layouts is what Snowflake's document processing capability exists for, and the documents can be processed where they are staged.

Retrieval over the account's own data combined with a generation function is the standard pattern for this, and it runs entirely inside the account.

Classifying free text against a set of categories is a standard Cortex task and is well suited to running close to the ticket data.

Question: 8

An airline's crew-rules assistant retrieves context from the Cortex Search service `crew_rules_svc`, defined with `TARGET_LAG = '1 hour'`, and generates answers with `AI_COMPLETE`. A rest-period rule was changed in the source table three days ago, yet the assistant still quotes the old rule, and `SEARCH_PREVIEW` returns the old text too. `DESCRIBE CORTEX SEARCH SERVICE crew_rules_svc` shows `INDEXING_STATE` as `SUSPENDED`.

What should the team do?

Response:

- A. Lower `TARGET_LAG` to one minute so refreshes pick up the rule change sooner
- B. Add an instruction to the `AI_COMPLETE` prompt telling the model to prefer the most recently dated version of any rule in the retrieved context
- C. Read `INDEXING_ERROR`, fix the failing source query, then run `ALTER CORTEX SEARCH SERVICE crew_rules_svc RESUME INDEXING`
- D. Run `ALTER CORTEX SEARCH SERVICE crew_rules_svc RESUME SERVING`, so queries reach the refreshed index

Answer: C

Explanation:

Indexing suspends itself after five consecutive refresh failures in the source query, and the error is reported in `INDEXING_ERROR`. Once the query is fixed, resuming indexing lets the refresh pick up the changed rule.

Question: 9

A travel booking site prototyped help-center search in a worksheet, calling `SNOWFLAKE.CORTEX.SEARCH_PREVIEW` against its Cortex Search service. The results and `@eq` filters

look right, and the site now needs exactly the same search behavior behind its customer-facing help widget, where response time matters.

How should the widget query the service?

Response:

- A. Call SEARCH_PREVIEW from a stored procedure that the widget invokes for each customer question
- B. Replace the service with a VECTOR_COSINE_SIMILARITY query over AI_EMBED output, avoiding the serving layer
- C. Send each request to the service's REST :query endpoint, or call search() through the Snowflake Python API, with the same query, columns and filter
- D. Load the questions into a table and run SEARCH_PREVIEW over its question column in one batch

Answer: C

Explanation:

The REST and Python APIs are the serving interfaces for applications and carry lower latency. They accept the same query, columns and filter object used in the prototype, so behavior carries over.

Question: 10

A model is invoked over a table that holds records for every region, and each regional team must see only its own rows in the results.

Which policy object enforces this?

Response:

- A. A masking policy on the sensitive columns
- B. A projection policy on the identifying columns
- C. An aggregation policy requiring a minimum group size
- D. A row access policy on the table

Answer: D

Explanation:

A row access policy filters which ROWS the invoking role can see, and it applies to the function's read exactly as it applies to any other query. Each team's calls therefore operate over its own slice without anything being duplicated per region.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>