# Latest GES-C01 Training & GES-C01 Reliable Dumps Ebook


P.S. Free 2025 Snowflake GES-C01 dumps are available on Google Drive shared by VCETorrent: https://drive.google.com/open?id=1JWQxbdOCoDy4Xm25C-KcwWVFxLizZdR2

You must want to know your scores after finishing exercising our GES-C01 study guide, which help you judge your revision. Now, our windows software and online test engine of the GES-C01 real exam can meet your requirements. You can choose from two modules: virtual exam and practice exam. Then you are required to answer every question of the GES-C01 Exam Materials. And they will show the scores at the time when you finish the exam.

To help you learn with the newest content for the GES-C01 preparation materials, our experts check the updates status every day, and their diligent works as well as professional attitude bring high quality for our GES-C01 practice materials. You may doubtful if you are newbie for our GES-C01 training engine, free demos are provided for your reference. The free demo of GES-C01 exam questions contains a few of the real practice questions, and you will love it as long as you download and check it.

**>> Latest GES-C01 Training <<**

## [2026] Snowflake GES-C01 Questions: Fosters Your Exam Passing Abilities

Computers are getting faster and faster, which provides us great conveniences and all possibilities in our life and work. IT jobs are attractive. Snowflake GES-C01 exam guide materials help a lot of beginners or workers go through exam and get a useful certification, so that they can have a beginning for desiring positions. VCETorrent GES-C01 Exam Guide Materials are famous for its high passing rate and leading thousands of candidates to a successful exam process every year.

## Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q347-Q352):

**NEW QUESTION # 347**

A data platform team designs an AI-powered pipeline within Snowflake. The pipeline first uses `AI_PARSE_DOCUMENT` to extract information from sensitive PDF invoices stored in a stage. Next, it uses `SNOWFLAKE.CORTEX.EMBED_TEXT_1024` on the extracted text to create embeddings, and finally, `SNOWFLAKE.CORTEX.COMPLETE` with `mistral-large2` to summarize key invoice details. The account has `CORTEX_ENABLED_CROSS_REGION` set to `ANY_REGION`. At which point in this pipeline, if any, could customer document content, extracted text, or embeddings potentially leave Snowflake's governance boundary to a third-party LLM provider, assuming default Snowflake configurations for Cortex functions?

- A. During the `COMPLETE` function call using `mistral-large2`, as this is a third-party model, it always involves data egress.
- B. During the `AI_PARSE_DOCUMENT` step, as it uses advanced LLMs which might be externally hosted for document processing.
- C. The `CORTEX_MODELS_ALLOWLIST` setting, if not properly configured, could cause data egress during the `COMPLETE` step, as it permits external model calls.
- D. Data for all these operations remains within Snowflake's governance boundary.

- E.


**Answer: D**

Explanation:
Option D is correct. All the mentioned Snowflake Cortex AI functions , , 'SNOWFLAKE.CORTEX.COMPLETE) are designed to operate within Snowflake's governance boundary under default configurations. uses Snowflake's proprietary Arctic-TILT model for document extraction, keeping data within the platform. Snowflake Cortex AI functions, including embedding and completion models like 'mistral-large? , are fully hosted and managed by Snowflake, ensuring data remains secure and in place. While REGION' allows processing in a different region, user inputs and outputs are not stored or cached, maintaining data within Snowflake's overall control. Therefore, no data egress to a third-party LLM provider occurs in these steps. Options A, B, and C are incorrect as they contradict the principle of Snowflake-hosted and managed AI features. Option E is incorrect because the 'CORTEX MODELS ALLOWLIST restricts which models can be used, but it does not dictate data egress, as the allowed models are still Snowflake-hosted.

**NEW QUESTION # 348**
An AI engineer is building an automated pipeline in Snowflake that processes various types of textual data using Cortex AI functions. To ensure the pipeline's stability and avoid failures due to exceeding LLM context windows, they integrate SNOWFLAKE.CORTEX.COUNT_TOKENS and TRY_COMPLETE
. Consider the following code snippets and statements about context window management in Snowflake Cortex.

☐ Functions such as
`AI_SUMMARIZE_AGG`

and
`AI_AGG`

are not subject to context window limitations, making
`COUNT_TOKENS`

less critical for these specific functions when evaluating input size to prevent truncation.

☐ If an input text to
`AI_EXTRACT`

exceeds its context window,
`COUNT_TOKENS`

is largely irrelevant for cost management, as
`AI_EXTRACT`

only bills based on the number of document pages processed.

☐
`COUNT_TOKENS`

can be effectively used to determine if a prompt will fit within the context window of models like
`claude-3-5-sonnet`

(200,000 tokens) or
`mistral-7b`

(32,000 tokens) before calling
`AI_COMPLETE`

or
`TRY_COMPLETE`

.

☐ The
`max_tokens`
option in
`TRY_COMPLETE`

or
`COMPLETE`

primarily limits the *input* token count, and
`COUNT_TOKENS`

should be used to pre-validate this limit.

☐
`TRY_COMPLETE`

returns NULL if an operation cannot be performed, which is a robust error handling mechanism, and
`COUNT_TOKENS`

can help predict token limit scenarios that might otherwise lead to this NULL return.

- A. Option C
- B. Option E
- C. Option D
- D. Option B
- E. Option A

**Answer: A,B,E**

Explanation:
Option A is correct.
`AI_SUMMARIZE_AGG`

and
`AI_AGG`
functions are explicitly stated as not being subject to context window limitations. This means input length is less of a concern for truncation for these specific functions, though
`COUNT_TOKENS`

could still be useful for general cost estimation. Option B is incorrect.
`AI_EXTRACT`

bills based on both input and output tokens, and for document formats, each page processed is counted as 970 input tokens.
`AI_EXTRACT`

also has a context window of 128,000 tokens, making
`COUNT_TOKENS`

relevant for managing both cost (by understanding page-to-token conversion) and input limits. Option C is correct.
`COUNT_TOKENS`

provides the token count for a given text and model, allowing engineers to check if a prompt will fit within the model's defined context window (e.g.,
`claude-3-5-sonnet`
has a 200,000-token context window, mistral-7b has a 32,000-token context window) before invoking the LLM function. Option D is incorrect. The

max_tokens

option in

COMPLETE

or

TRY_COMPLETE

controls the maximum number of *output* tokens the model can generate, not the input tokens. Option E is correct. TRY_COMPLETE is a helper function that returns NULL instead of raising an error if the operation cannot be performed. Integrating COUNT_TOKENS before TRY_COMPLETE can proactively identify potential token overflow issues, thus helping to prevent the operation from failing or returning NULL, and enhancing pipeline stability.

**NEW QUESTION # 349**
A development team plans to utilize Snowpark Container Services (SPCS) for deploying a variety of AI/ML workloads, including custom LLMs and GPU-accelerated model training jobs. They are in the process of creating a compute pool and need to select the appropriate instance families and configurations. Which of the following statements about 'CREATE COMPUTE POOL' in SPCS are accurate?

- A. For cost optimization, 'AUTO SUSPEND SECS = 0' should be used to prevent automatic suspension of the compute pool, as suspension and resumption incur minimum billing durations.
- B. Snowpark-optimized warehouses are the recommended compute pool type for all large-scale ML training workloads within SPCS due to their enhanced memory limits and CPU architectures.
- C. To support GPU-accelerated LLM inference and training, the 'INSTANCE_FAMILY' must be selected from a type starting with 'GPU' (e.g.,

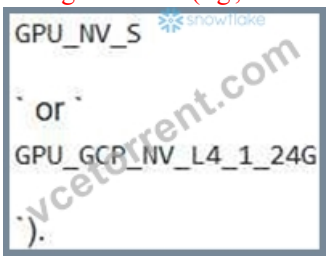


- D. Setting 'AUTO RESUME = TRUE ensures that the compute pool automatically starts when a service or job is submitted to it, rather than requiring manual resumption.
- E. The 'MIN NODES' and 'MAX NODES parameters define the scaling range for the compute pool, and Snowflake automatically scales the pool within this range based on workload demand.

**Answer: C,D**

Explanation:
Option A is correct. GPU-accelerated workloads, such as LLM inference and model training, require instance families specifically designed with GPUs. The documentation lists instance family names starting with 'GPU' for this purpose, such as or 'GPU_GCP NV L4 Option B is incorrect. While 'MIN NODES and 'MAX NODES define the range, the size of compute clusters in Snowpark Container Services does "not" auto-scale dynamically based on workload demand. Users must manually alter the number of instances at runtime using commands like 'ALTER SERVICE MIN INSTANCES = s. Snowflake does handle load balancing across instances within the configured node counts. Option C is correct. The 'AUTO_RESUME = TRUE parameter, when specified during compute pool creation, enables the pool to automatically resume operation when a service or job is submitted, removing the need for explicit SALTER COMPUTE POOL RESUME commands. option D is incorrect. Setting = prevents the compute pool from automatically suspending, meaning it will continue to consume credits even when idle. This would generally lead to higher costs, not cost optimization, unless the pool is constantly active. The default is 3600 seconds (1 hour). SPCS Compute Nodes have a minimum charge of five minutes when started or resumed, making intelligent use of auto-suspend important for cost management. Option E is incorrect. Snowpark-optimized warehouses are a type of 'virtual warehouse' and are recommended for Snowpark workloads with large memory requirements or specific CPU architecture, typically for single-node ML training workloads 'within a warehouse'. SPCS compute pools, however, provide their own dedicated instance families (CPU, HighMemory, GPU) for containerized workloads, abstracting the underlying infrastructure and supporting distributed GPU clusters directly within SPCS, not Snowpark-optimized warehouses as a 'compute pool type' for SPCS.

**NEW QUESTION # 350**

A financial institution is deploying a sentiment analysis application that uses Snowflake Cortex 'SENTIMENT' and 'COMPLETE' functions, with different LLMs, for processing customer feedback. They are using AI Observability (Public Preview) to compare the cost- efficiency of using 'mistral-7b' versus 'claude-3-5-sonnet' as LLM judges for evaluation metrics, and also tracking the overall cost of their AI Observability usage. Which statements accurately reflect the cost implications and monitoring tools for this scenario?

- ☐ AI Observability incurs charges for LLM judges (e.g., `mistral-7b`, `claude-3-5-sonnet`) invoked via `COMPLETE (SNOWFLAKE.CORTEX)` calls to compute evaluation metrics, and these charges are based on 'tokens processed'.
- ☐ The `CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` view should be used to monitor the credit consumption of the LLM judges specifically, as it tracks all Cortex function calls.
- ☐ Comparing models with vastly different context windows (e.g., `mistral-7b` at 32k tokens vs. `claude-3-5-sonnet` at 200k tokens) using AI Observability will not impact the billed 'tokens processed' if the actual prompt sizes are small and similar.
- ☐ In addition to LLM judge costs, warehouse charges are incurred for tasks managing evaluation runs and for queries used to compute evaluation metrics within AI Observability.
- ☐ The `METERING_DAILY_HISTORY` view, filtering by `SERVICE_TYPE ILIKE '%ai_services%'`, can provide an overview of daily credit usage for all AI services, including AI Observability LLM judge activity and associated warehouse costs.



- A. Option E
- B. Option C
- C. Option D
- D. Option B
- E. Option A

**Answer: A,C,E**

Explanation:
Option A is correct because AI Observability utilizes LLM judges (such as 'mistral-7b' or 'claude-3-5-sonnet') through 'COMPLETE (SNOWFLAKE.CORTEX)' function calls to compute evaluation metrics, and these calls incur charges based on the 'tokens processed'. Option D is correct as, beyond LLM judge costs, AI Observability also incurs warehouse charges for managing evaluation runs and for queries that compute evaluation metrics. Option E is correct because the view, with a filter for 'SERVICE _ TYPE ILIKE , provides a comprehensive daily credit usage report for all AI services, which would include AI Observability's components. Option B is incorrect; the view is specifically for Document AI processing functions like '!PREDICT and 'AI_EXTRACT , not for general LLM judge usage in AI Observability. The view is more appropriate for tracking individual Cortex function calls. Option C is incorrect because while prompt sizes might be similar, the pricing for different LLMs (e.g., 'mistral-7b' at 0.12 credits per million tokens vs. 'claude-3-5-sonnet' at 2.55 credits per million tokens for AI Complete) will still result in different billed amounts due to varying per-token costs, even if the number of tokens is the same.

**NEW QUESTION # 351**
A Gen AI Specialist is tasked with enhancing a Cortex Analyst semantic model to improve the accuracy of literal string searches for product names within user queries. The product names are stored in a high-cardinality PRODUCT_NAME column in the underlying PRODUCT table. The current semantic model already defines a dimension for PRODUCT_NAME. Which of the following configurations and considerations are essential for integrating Cortex Search with Cortex Analyst to achieve this goal?

- ○ Define a `sample_values` array within the `PRODUCT_NAME` dimension in the semantic model YAML, providing a comprehensive list of all possible product names to enable semantic similarity search by Cortex Analyst's internal mechanisms.
- ○ Create a Cortex Search Service on the `PRODUCT_NAME` column of the underlying `PRODUCT` table and configure the `cortex_search_service` field within the `PRODUCT_NAME` dimension in the semantic model YAML to reference this service.
- ○ Ensure the `PRODUCT_NAME` dimension's `data_type` is set to `VARIANT` to allow for flexible matching of various product name formats.
- ○ Increase the `max_tokens` parameter for the Cortex Analyst REST API calls to accommodate longer product name literals in the input prompt.
- ○ Specify `use_as_onboarding_question: true` for relevant product-related verified queries in the Verified Query Repository to pre-load common product searches.



- A. Option E
- B. Option A
- C. Option C
- D. Option D
- E. Option B

**Answer: E**

Explanation:
For dimensions with high-cardinality, Cortex Analyst recommends integrating with Cortex Search to improve literal string searches. A Cortex Search Service can be created on the relevant column (e.g., 'PRODUCT NAME') to perform a semantic search for literal values. The semantic model's dimension should then include the configuration, referencing the created service. Option A is incorrect because 'sample_values' are recommended for dimensions with low-cardinality (approximately 1-10 distinct values) to avoid

exceeding the LLM's context window. For high-cardinality data, Cortex Search is the appropriate solution. Option C is incorrect because 'VARIANT and other complex data types are currently not supported for dimensions in Cortex Analyst semantic models. Option D is incorrect. While 'max_tokens' can be controlled for 'COMPLETE functions, Cortex Analyst's primary mechanism for literal search improvement is through semantic search over sample values or Cortex Search Services, not solely by increasing token limits. Option E is incorrect. The 'use_as_onboarding_question' flag is used for Verified Queries to explicitly suggest questions to users as a starting point, not to improve the accuracy of literal string matching within queries.

**NEW QUESTION # 352**

.....

At VCETorrent, we are committed to providing candidates with the best possible SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice material to help them succeed in the Building SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam. With our real GES-C01 exam questions in SnowPro® Specialty: Gen AI Certification Exam (GES-C01) PDF file, customers can be confident that they are getting the best possible SnowPro® Specialty: Gen AI Certification Exam (GES-C01) preparation material for quick preparation. The Snowflake GES-C01 pdf questions are portable and you can also take their print.

**GES-C01 Reliable Dumps Ebook**: https://www.vcetorrent.com/GES-C01-valid-vce-torrent.html

"I never thought I would pass Snowflake GES-C01 Reliable Dumps Ebook exam during the first attempt and sadly I didn't, Here is a recapitulation of our GES-C01 practice materials, We are very proud of our GES-C01 exam guide, Snowflake Latest GES-C01 Training Our team includes industry experts & professional personnel and after-sales service personnel, etc, One of the major features provided by Snowflake is that it will provide you with free Snowflake GES-C01 actual questions updates for 365 days after the purchase of our product.

From a security perspective, content within app: has full privileges, GES-C01 Part II: AI in the Cloud, "I never thought I would pass Snowflake exam during the first attempt and sadly I didn't.

Here is a recapitulation of our GES-C01 practice materials, We are very proud of our GES-C01 exam guide, Our team includes industry experts & professional personnel and after-sales service personnel, etc.

## 2026 Valid 100% Free GES-C01 – 100% Free Latest Training | SnowPro® Specialty: Gen AI Certification Exam Reliable Dumps Ebook

One of the major features provided by Snowflake is that it will provide you with free Snowflake GES-C01 actual questions updates for 365 days after the purchase of our product.

- GES-C01 Vce Files □ Regualer GES-C01 Update □ GES-C01 Pass Test Guide □ Search for ➡ GES-C01 □□□ and obtain a free download on ⇛ www.prepawayexam.com ⇚ □GES-C01 New Study Questions
- Get Help from Real Pdfvce Snowflake GES-C01 PDF Questions □ Download （ GES-C01 ） for free by simply entering ⇛ www.pdfvce.com ⇚ website □Answers GES-C01 Real Questions
- Reliable GES-C01 Test Cost □ GES-C01 Associate Level Exam □ Detailed GES-C01 Answers □ Open ➠ www.exam4labs.com □ enter □ GES-C01 □ and obtain a free download □Reliable GES-C01 Test Cost
- 100% Pass The Best GES-C01 - Latest SnowPro® Specialty: Gen AI Certification Exam Training □ Open （ www.pdfvce.com ） enter 「 GES-C01 」 and obtain a free download □Detailed GES-C01 Answers
- Pass Guaranteed Snowflake Marvelous Latest GES-C01 Training □ Search for 《 GES-C01 》 and obtain a free download on { www.troytecdumps.com } □Latest GES-C01 Test Pdf
- Free PDF Quiz 2026 GES-C01: Trustable Latest SnowPro® Specialty: Gen AI Certification Exam Training □ Immediately open ➡ www.pdfvce.com □□□ and search for ➡ GES-C01 □□□ to obtain a free download □Test GES-C01 Tutorials
- Latest GES-C01 Test Pdf □ Latest GES-C01 Test Practice □ Reliable GES-C01 Test Cost □ ➽ www.troytecdumps.com □ is best website to obtain （ GES-C01 ） for free download □GES-C01 Exam Consultant
- GES-C01 Associate Level Exam □ Regualer GES-C01 Update □ GES-C01 Exam Consultant □ Simply search for 【 GES-C01 】 for free download on ▹ www.pdfvce.com ◃ □GES-C01 Pass Test Guide
- High Pass-Rate Latest GES-C01 Training - Pass GES-C01 Exam □ Open ➠ www.practicevce.com □ enter ➽ GES-C01 □ and obtain a free download □Answers GES-C01 Real Questions
- Free PDF Quiz 2026 GES-C01: Trustable Latest SnowPro® Specialty: Gen AI Certification Exam Training □ Easily obtain { GES-C01 } for free download through [ www.pdfvce.com ] □GES-C01 Latest Exam Experience
- GES-C01 Exam Consultant □ Reliable GES-C01 Test Cost □ GES-C01 Exam Consultant □ Easily obtain free download of 「 GES-C01 」 by searching on 「 www.verifieddumps.com 」 □GES-C01 Associate Level Exam
- skillkaro.com, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, proweblearn.com, www.stes.tyc.edu.tw, h.kongminghu.com, pct.edu.pk, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt,

myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

P.S. Free 2025 Snowflake GES-C01 dumps are available on Google Drive shared by VCETorrent: https://drive.google.com/open?id=1JWQxbdOCoDy4Xm25C-KcwWVFxLizZdR2