Study GES-C01 Test - Valid GES-C01 Test Cram

BONUS!!! Download part of ActualtestPDF GES-C01 dumps for free: https://drive.google.com/open?id=1ZyMQeGJJi50eIqUmmcli9DxaDudqx6Z7

If you find any quality problems of our GES-C01 or you do not pass the exam, we will unconditionally full refund. ActualtestPDF is professional site that providing Snowflake GES-C01 Questions and answers, it covers almost the GES-C01 full knowledge points.

Are you still feeling uncomfortable about giving up a lot of time to entertain, work or accompany your family and friends in preparation for the exam? Using GES-C01 quiz torrent, you can spend less time and effort reviewing and preparing, which will help you save a lot of time and energy. Whether you are a worker or student, you will save much time to do something whatever you want. It only needs 5-10 minutes after you pay for our GES-C01 learn torrent that you can learn it to prepare for your exam. Actually, if you can guarantee that your effective learning time with GES-C01 test preps are up to 20-30 hours, you can pass the exam.

>> Study GES-C01 Test <<

Specifications of ActualtestPDF Snowflake GES-C01 Exam Preparation Material

The students can give unlimited to track the performance of their last given tests in order to see their mistakes and try to avoid them while giving the final test. Customers of ActualtestPDF will receive updates till 1 year after their purchase. Anyone can try a free demo of the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice material before making purchase. There is a 24/7 available support system that assists users whenever they are stuck in any problem or issues. This product is a complete package and a blessing for those who want to pass the Snowflake GES-C01 test in a single try.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q95-Q100):

NEW QUESTION # 95
A data scientist is tasked with improving the accuracy of an LLM-powered chatbot that answers user questions based on internal company documents stored in Snowflake. They decide to implement a Retrieval Augmented Generation (RAG) architecture using Snowflake Cortex Search. Which of the following statements correctly describe the features and considerations when leveraging Snowflake Cortex Search for this RAG application?

- A. The
- B. Cortex Search automatically handles text chunking and embedding generation for the source data, eliminating the need for manual ETL processes for these steps.
- C. For optimal search results with Cortex Search, source text should be pre-split into chunks of no more than 512 tokens, even when using models with larger context windows like
  `snowflake-arctic-embed-l-v2.0-8k`
- D. To create a Cortex Search Service, one must explicitly specify an embedding model and manually manage its underlying

infrastructure, similar to deploying a custom model via Snowpark Container Services.
- E. Enabling change tracking on the source table for the Cortex Search Service is optional; the service will still refresh automatically even if change tracking is disabled.

Answer: A,B,C

Explanation:
Option A is correct because Cortex Search is a fully managed service that gets users started with a hybrid (vector and keyword) search engine on text data in minutes, without needing to worry about embedding, infrastructure maintenance, or index refreshes. Option B is incorrect because Cortex Search is a fully managed service; users do not need to manually manage the embedding model infrastructure. A default embedding model is used if not specified. Option C is correct because, for best search results with Cortex Search, Snowflake recommends splitting text into chunks of no more than 512 tokens, as smaller chunks typically lead to higher retrieval and downstream LLM response quality, even with models that have larger context windows. Option D is correct because the SNOWFLAKE.CORTEX.SEARCH_PREVIEW' function allows users to test the search service to confirm it is populated with data and serving reasonable results for a given query. Option E is incorrect because change tracking is required on the source table for the Cortex Search Service to function correctly and reflect updates to the base data.

NEW QUESTION # 96
A machine learning team is leveraging the Snowflake Model Registry to manage diverse models, including a custom Python utility for data preprocessing that they wish to make available as a model method. Which of the following statements accurately describe capabilities or considerations when logging models and their associated artifacts and methods in the Model Registry?

- The Snowflake Model Registry supports built-in types such as Scikit-learn, XGBoost, and PyTorch, but does not allow logging custom Python objects or processing code directly as models.
- To include additional local files, such as configuration files or custom scripts, with a logged model, the `user_files` argument must be used in `log_model`, mapping stage subdirectories to local file paths.
- Once a model version is logged, its methods can be invoked using either `mv.run()` in Python or through service functions named `<service_name>!<method_name>` in SQL, after the model has been deployed to SPCS.
- The `function_type` option within `method_options` in `log_model` allows specifying whether a model method should be exposed as a `FUNCTION` or `TABLE_FUNCTION` in SQL, influencing how data is processed.
- The maximum total model size for models deployed to a Snowflake warehouse is 5 GB, whereas models deployed to SPCS have no such size limitations.

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

Answer: A,C,D,E

Explanation:
Option B is correct because the `user_files` argument in `log_model` is specifically designed for including additional local files (e.g., configuration, scripts) with the model, by mapping them to stage subdirectories. Option C is correct because after a model is deployed to SPCS, Snowflake Model Serving creates service functions named `<service_name>!<method_name>` that allow invoking the model's methods from SQL. Model methods can also be called via `mv.run()` in Python. Option D is correct because the `function_type` option within `method_options` allows developers to specify if a model method should be registered as a `FUNCTION` or `TABLE_FUNCTION` when exposed in SQL, affecting its input/output processing. Option E is correct. The maximum total model size for warehouse-deployed models is 5 GB. For models running on Snowpark Container Services, these size restrictions are eased or eliminated, allowing for much larger models. Option A is incorrect; the Model Registry is flexible enough to support not only built-in types (like Scikit-learn, XGBoost, PyTorch) but also custom processing code and previously-trained models.

NEW QUESTION # 97
A global enterprise has Snowflake accounts in various regions, including a US East (Ohio) account where a critical application is deployed. They need to use AI_COMPLETE with the claude-3-5-sonnet model for real-time customer support, but this model is not natively available in US East (Ohio) for direct AI_COMPLETE usage. The Snowflake administrator considers enabling cross-region inference. Which statements accurately reflect the considerations and characteristics of cross-region inference in Snowflake Cortex?

- A. Setting the CORTEX_ENABLED_CROSS_REGION account parameter to 'ANY_REGION' in the US East (Ohio) account would enable inference requests for claude-3-5- sonnet to be processed in any region where it is natively available.
- B. The CORTEX_ENABLED_CROSS_REGION parameter can be configured at the session level to temporarily enable

cross-region inference for specific workloads.

- C. Cross-region inference automatically caches user inputs and generated outputs to reduce latency for subsequent requests to the same model.
- D. Latency between regions for cross-region inference is negligible and consistently low, irrespective of cloud provider infrastructure.
- E. Cross-region inference is not supported in U.S. SnowGov regions for either inbound or outbound inference requests.

Answer: A,E

Explanation:
Option B is correct because setting the parameter to 'ANY_REGION' enables inference requests to be CORTEX_ENABLED_CROSS_REGION processed in a different region from the default, thereby allowing access to models not natively supported in the local region. For example, claude- is 3-5- sonnet available in AWS US East 1 (N. Virginia), which could be accessed from US East (Ohio) via cross-region inference. Option C is 3-5 -sonnet correct as cross-region inference is explicitly not supported in U.S. SnowGov regions. Option A is incorrect because user inputs, service generated prompts, and outputs are not stored or cached during cross-region inference. Option D is incorrect; latency depends on the cloud provider infrastructure and network status, and testing is recommended. Option E is incorrect because CORTEX_ENABLED_CROSS_REGION is an account-level parameter, not a session parameter.

NEW QUESTION # 98
A Snowflake account administrator in an Azure East US 2 region needs to enable users to access a new, highly capable LLM, 'claude-3-5-sonnet' , which is currently only natively available in AWS regions via Snowflake Cortex. The administrator also wants to ensure that only specific, approved LLMs can be used across the organization. Which configuration steps are necessary for the administrator to achieve these requirements?

- A. Since 'claude-3-5-sonnet' is an OpenAl model, the administrator must enable the 'ENABLE_CORTEX ANALYST MODEL AZURE OPENAI' account parameter, and then the model will automatically be available for cross-region inference without further action.
- B. Create a 'COMPUTE POOL' with a 'GPU NV_S instance family in Azure East US 2, and then deploy a custom PyCaret model of 'claude-3-5-sonnet' to this pool through the Snowflake Model Registry.
- C. Set the account parameter to include 'claude-3-5-sonnet' , and then set the account parameter to 'TRUE to allow cross-region inference for all Cortex features.
- D. The 'CORTEX ENABLED CROSS REGION' parameter allows access to models in other regions, but access to specific LLMs is controlled solely by individual user privileges granted directly on the model objects, not by an account-level allowlist.
- E. Grant the 'SNOWFLAKE.CORTEX_USER database role to the relevant user roles. Set the account parameter to 'ANY REGION' or a list including an AWS region where 'claude-3-5-sonnet' is natively available. Additionally, configure the 'CORTEX MODELS ALLOWLIST to explicitly permit 'claude-3-5-sonnet' and other desired models.

Answer: E

Explanation:
Option A is partially correct but incomplete. Setting to 'TRUE or 'ANY_REGION' allows cross-region inference. However, ' TRUE' is not a valid value for 'CORTEX ENABLED CROSS REGION in the provided sources, it can be a list of regions or 'ANY_REGION'. The does control which models can be used. Option B is correct. To call Snowflake Cortex AI functions, the 'SNOWFLAKE.CORTEX USER database role is required. To access 'claude-3-5-sonnet' (which is available via cross-cloud inference) from an Azure region when it's natively available in an AWS region, the 'CORTEX ENABLED CROSS REGION' parameter must be configured to allow it, either by specifying 'ANY REGION' or listing the relevant AWS region. Additionally, the parameter is used by administrators to restrict or allow access to specific LLMs within Snowflake. 'claude-3-5-sonnet' is a supported model for 'AI_COMPLETE' and 'COMPLETE. Option C is incorrect. 'claude-3-5-sonnet' is an Anthropic model, not an OpenAl GPT model. Furthermore, is a legacy parameter for Cortex Analyst specifically for Azure OpenAl models, and its use is discouraged. Option D is incorrect. While 'COMPUTE POOL' and instance families are used for Snowpark Container Services and Hugging Face models can be deployed this way, this scenario involves directly using a Snowflake-hosted Cortex LLM ('claude-3-5-sonnet') rather than deploying a custom external model, and 'claude-3-5-sonnet' is not a PyCaret model. Option E is incorrect. While RBAC applies to model objects, the 'CORTEX MODELS ALLOWLIST' is an account-level parameter used by administrators to limit which LLMs can be used, overriding or complementing individual object grants.

NEW QUESTION # 99
A multi-national corporation uses Snowflake across several AWS regions. Their primary operational Snowflake account is in AWS US East (Ohio), but they need to leverage a specific AI_COMPLETE model, llama4-maverick, which is natively available in AWS

US East 1 (N. Virginia) but not in US East (Ohio). To address this, the Snowflake administrator enables cross-region inference for their US East (Ohio) account.

- A. The llama4-maverick model is listed as natively available in AWS US East 1 (N. Virginia) and is supported for cross-region inference (AWS US Cross-Region), validating it as a suitable target for inference from US East (Ohio).
- B. Cross-region inference is fully supported for AI_COMPLETE in U.S. SnowGov regions for both inbound and outbound inference requests, provided the target model is available in the respective SnowGov region.
- C. User inputs, service-generated prompts, and the generated outputs from cross-region AI_COMPLETE calls are automatically stored or cached in the remote processing region to optimize performance for subsequent identical requests.
- D. To enable cross-region inference for the US East (Ohio) account, the administrator would execute the command: ALTER ACCOUNT SET AWS_US' ; to allow inference requests to be processed in any AWS US region where the model is available. CORTEX_ENABLED_CROSS REGION =AWS_US' ;
- E. The query latency for cross-region inference with AI_COMPLETE is consistently low and predictable, as Snowflake's architecture is designed to completely negate the impact of geographical distance and network variations.

Answer: A,D

Explanation:
Option A is correct because the CORTEX_ENABLED_CROSS_REGION account parameter is used to enable cross-region inference. Setting it CORTEX_ENABLED_CROSS_REGION would permit inference requests to be processed in any AWS US region, such as N. Virginia, from a local AWS US region like Ohio. to 'Aws us' Option B is incorrect because user inputs, service generated prompts, and outputs are explicitly not stored or cached during cross-region inference. Option C is incorrect as cross-region inference is not supported in U.S. SnowGov regions for either inbound or outbound inference requests. Option D is correct because the sources indicate that AI COMPLETE (llama4-maverick) is natively available in AWS US East 1 (N. Virginia) and is supported for cross-region inference within AWS US regions, making it a valid target for the account in US East (Ohio). Option E is incorrect because latency between regions depends on the cloud provider infrastructure and network status, and Snowflake recommends testing specific use cases with cross-region inference enabled.

NEW QUESTION # 100
.....

Registration

- Test GES-C01 Free □ GES-C01 New Dumps Files □ Test GES-C01 Free □ Search for ➥ GES-C01 □ on ➡ www.exam4labs.com □□□ immediately to obtain a free download □GES-C01 Exam Price
- GES-C01 Exam Price ➽ GES-C01 Exam Price □ GES-C01 Exam Price □ Download { GES-C01 } for free by simply entering { www.pdfvce.com } website □GES-C01 Practice Test Pdf
- GES-C01 Test Questions □ GES-C01 New Dumps Files □ Test GES-C01 Free □ Search for ▸ GES-C01 ◂ on （ www.examdiscuss.com ） immediately to obtain a free download □Exam GES-C01 Learning
- GES-C01 Exam Price □ Valid GES-C01 Exam Papers □ GES-C01 Test Questions □ The page for free download of 《 GES-C01 》 on □ www.pdfvce.com □ will open immediately □Exam GES-C01 Learning
- GES-C01 Certification Test Answers □ Valid GES-C01 Test Registration □ GES-C01 Test Questions □ The page for free download of 「 GES-C01 」 on 【 www.examdiscuss.com 】 will open immediately □Real GES-C01 Dumps Free
- Test GES-C01 Simulator □ GES-C01 Discount □ Test GES-C01 Free □ Open （ www.pdfvce.com ） enter ➠ GES-C01 □ and obtain a free download □GES-C01 Excellect Pass Rate
- GES-C01 Discount □ Exam GES-C01 Guide □ GES-C01 Exam Price ♣ Enter □ www.prepawaypdf.com □ and search for ➠ GES-C01 □ to download for free □Real GES-C01 Dumps Free
- 2026 Study GES-C01 Test | High Pass-Rate 100% Free Valid SnowPro® Specialty: Gen AI Certification Exam Test Cram □ Search for ▸ GES-C01 ◂ and download it for free immediately on ➽ www.pdfvce.com □ □Learning GES-C01 Mode
- GES-C01 Free Download Demo - GES-C01 Latest Exam Tutorial - GES-C01 Valid Study Reviews □ Search on ▸ www.prep4sures.top ◂ for 【 GES-C01 】 to obtain exam materials for free download □GES-C01 Practice Test Pdf
- minaegkt497378.blogginaway.com, luluqlie490776.sasugawiki.com, arunipae665373.onzeblog.com, tedktui481891.daneblogger.com, www.stes.tyc.edu.tw, amaanphun839960.ktwiki.com, socialclubfm.com, brendaprnh368109.theobloggers.com, iwanshfa805146.gynoblog.com, heathaskm167550.fliplife-wiki.com, Disposable vapes

P.S. Free 2026 Snowflake GES-C01 dumps are available on Google Drive shared by ActualtestPDF: https://drive.google.com/open?id=1ZyMQeGJJi50eIqUmmcli9DxaDudqx6Z7