

Updated Snowflake DSA-C03 Testkings, DSA-C03 Valid Exam Pattern



BTW, DOWNLOAD part of Pass4suresVCE DSA-C03 dumps from Cloud Storage: <https://drive.google.com/open?id=1ettSqdk9kc1SA-QWfqNxxwfy-IUv6XMhW>

With rigorous analysis and summary of DSA-C03 exam, we have made the learning content easy to grasp and simplified some parts that beyond candidates' understanding. In addition, we add diagrams and examples to display an explanation in order to make the interface more intuitive. Our DSA-C03 exam questions will ease your pressure of learning, using less Q&A to convey more important information, thus giving you the top-notch using experience if you study with our DSA-C03 Training Materials. And with the high pass rate of 99% to 100%, the DSA-C03 exam will be a piece of cake for you.

If you are preparing for the DSA-C03 Questions and answers, and like to practice it in your spare time, then you should consider the DSA-C03 exam dumps of our company. DSA-C03 Online test engine is convenient and easy to study, it supports all web browsers. Besides you can practice online anytime. With all the benefits like this, you can choose us bravely. With this version, you can pass the exam easily, and you don't need to spend the specific time for practicing, just your free time is ok.

>> Updated Snowflake DSA-C03 Testkings <<

DSA-C03 Valid Exam Pattern | Reliable DSA-C03 Real Test

If you have questions about us, you can contact with us at any time via email or online service. We will give you the best suggestions on the DSA-C03 study guide. And you should also trust the official cDSA-C03 certification. Or, you can try it by yourself by free downloading the demos of the DSA-C03 learning braindumps. I believe you will make your own judgment. We are very confident in our DSA-C03 exam questions.

Snowflake SnowPro Advanced: Data Scientist Certification Exam Sample Questions (Q222-Q227):

NEW QUESTION # 222

You are building a data science pipeline in Snowflake to predict customer churn. The pipeline includes a Python UDF that uses a pre-trained scikit-learn model stored as a binary file in a Snowflake stage. The UDF needs to load this model for prediction. You've

encountered an issue where the UDF intermittently fails, seemingly related to resource limits when multiple concurrent queries invoke the UDF. Which of the following strategies would best optimize the UDF for concurrency and resource efficiency, minimizing the risk of failure?

- A. Load the scikit-learn model inside the UDF function on every invocation to ensure the latest version is used.
- B. Increase the memory allocated to the Snowflake warehouse to accommodate multiple UDF invocations.
- **C. Implement a global, lazy-loaded cache for the scikit-learn model within the UDF's module. The model is loaded only once during the first invocation and shared across subsequent calls. Protect the loading process with a lock to prevent race conditions in concurrent environments.**
- D. Utilize Snowflake's session-level caching by storing the loaded model in 'session.get('model')' to be reused across multiple UDF calls within the same session. Reload the model if 'session.get('model')' is None.
- E. Load the scikit-learn model outside the UDF function in the global scope of the module so that all invocations share the same loaded model instance. Use the 'context.getExecutionContext()' to track execution, making sure it is thread safe.

Answer: C

Explanation:

Option D provides the most efficient and robust solution. Loading the model only once (lazy loading) reduces overhead. A global cache ensures reusability. A lock is crucial to prevent race conditions during the initial loading in a concurrent environment. Option A is inefficient due to repeated loading. Option B is problematic because Snowflake UDFs do not directly support global variables in a thread-safe manner. Option C is incorrect as 'session.get' is not a valid Snowflake API for Python UDFs and lacks thread safety. Option E, while potentially helpful, doesn't address the underlying inefficiency of repeatedly loading the model.

NEW QUESTION # 223

You are building a model training pipeline in Snowflake using Snowpark Python. You want to leverage a pre-trained model from Hugging Face Transformers for a text classification task, fine-tuning it with your own labeled data stored in a Snowflake table named 'training_data'. You've chosen the 'transformers' library and plan to use a 'transformers.pipeline' for inference. Which of the following code snippets, when integrated into your Snowpark Python application, will correctly download the pre-trained model and tokenizer, prepare the data, perform fine-tuning, and then save the fine-tuned model to a Snowflake stage?

- ```
from snowflake.snowpark import Session import transformers def train_model(session: Session, model_name: str, stage_name: str, table_name: str): model = transformers.AutoModelForSequenceClassification.from_pretrained(model_name) tokenizer = transformers.AutoTokenizer.from_pretrained(model_name) training_data = session.table(table_name).to_pandas() trainer = transformers.Trainer(model=model, train_dataset=training_data, tokenizer=tokenizer) trainer.train() model.save_pretrained(f'@{stage_name}') tokenizer.save_pretrained(f'@{stage_name}')
```
- ```
from snowflake.snowpark import Session from transformers import AutoModelForSequenceClassification, AutoTokenizer, Trainer, TrainingArguments def train_model(session: Session, model_name: str, stage_name: str, table_name: str): model = AutoModelForSequenceClassification.from_pretrained(model_name) tokenizer = AutoTokenizer.from_pretrained(model_name) training_data = session.table(table_name).to_pandas() training_args = TrainingArguments(output_dir='tmp_results', evaluation_strategy='epoch') trainer = Trainer(model=model, args=training_args, train_dataset=training_data.to_pandas(), tokenizer=tokenizer) trainer.train() trainer.save_model(f'@{stage_name}')
```
- ```
from snowflake.snowpark import Session import torch from transformers import AutoModelForSequenceClassification, AutoTokenizer, Trainer, TrainingArguments, TextClassificationPipeline def train_model(session: Session, model_name: str, stage_name: str, table_name: str): model = AutoModelForSequenceClassification.from_pretrained(model_name) tokenizer = AutoTokenizer.from_pretrained(model_name) training_data = session.table(table_name).to_pandas() training_args = TrainingArguments(output_dir='tmp_results', evaluation_strategy='epoch') trainer = Trainer(model=model, args=training_args, train_dataset=training_data, tokenizer=tokenizer) trainer.train() trainer.save_model(f'@{stage_name}') # Create a pipeline for inference and upload it to the stage pipeline = TextClassificationPipeline(model=model, tokenizer=tokenizer, device=torch.device('cuda' if torch.cuda.is_available() else 'cpu')) session.add_dependency(f'@{stage_name}')
```
- ```
from snowflake.snowpark import Session from transformers import AutoModelForSequenceClassification, AutoTokenizer, Trainer, TrainingArguments def train_model(session: Session, model_name: str, stage_name: str, table_name: str): model = AutoModelForSequenceClassification.from_pretrained(model_name) tokenizer = AutoTokenizer.from_pretrained(model_name) training_data = session.table(table_name).to_pandas() training_args = TrainingArguments(output_dir='/tmp', evaluation_strategy='epoch') trainer = Trainer(model=model, args=training_args, train_dataset=training_data, tokenizer=tokenizer) trainer.train() trainer.save_model(f'@{stage_name}')
```
- ```
from snowflake.snowpark import Session from transformers import AutoModelForSequenceClassification, AutoTokenizer, Trainer, TrainingArguments def train_model(session: Session, model_name: str, stage_name: str, table_name: str): model = AutoModelForSequenceClassification.from_pretrained(model_name) tokenizer = AutoTokenizer.from_pretrained(model_name) training_data = session.table(table_name).to_pandas() training_args = TrainingArguments(output_dir='./results', evaluation_strategy='epoch') trainer = Trainer(model=model, args=training_args, train_dataset=training_data, tokenizer=tokenizer) trainer.train() trainer.save_model(f'@{stage_name}')
```

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

**Answer: D**

Explanation:

The correct answer is B. It correctly uses the 'transformers' library with Snowpark Python. It downloads the model and tokenizer using 'AutoTokenizer.from\_pretrained'. TrainingArguments are configured with output\_dir and evaluation\_strategy. It reads training data using session.table. Trainer properly configured and finally Trainer saves the trained model in specified 'stage\_name'. Option A is incorrect because it missing 'TrainingArguments' configuration and uses general function, which may not be optimal for the Trainer setup. Option C is incorrect because incorrect use case. Option D and E is incorrect because 'TrainingArguments' output\_dir is local folder that cannot be written by Trainer.

### NEW QUESTION # 224

Consider you are working on a credit risk scoring model using Snowflake. You have a table 'credit\_data' with the following schema: 'customer\_id', 'age', 'income', 'credit\_score', 'loan\_amount', 'loan\_duration', 'defaulted'. You want to create several new features using Snowflake SQL to improve your model. Which combination of the following SQL statements will successfully create features for age groups, income-to-loan ratio, and interaction between credit score and loan amount using SQL in Snowflake? Choose all that apply.

```
CREATE OR REPLACE TEMPORARY TABLE credit_data_enriched AS
SELECT
 ,
 income / loan_amount AS income_to_loan_ratio
FROM credit_data;
```

- A.
- B.

```
ALTER TABLE credit_data ADD COLUMN credit_score_loan_interaction FLOAT;
UPDATE credit_data SET credit_score_loan_interaction = credit_score * loan_amount;
```

```
ALTER TABLE credit_data ADD COLUMN age_group VARCHAR;
UPDATE credit_data SET age_group = CASE
 WHEN age < 30 THEN 'Young'
 WHEN age BETWEEN 30 AND 50 THEN 'Adult'
 ELSE 'Senior'
END;
```

- C.

```
CREATE OR REPLACE TABLE credit_data_transformed AS
SELECT customer_id,
 age,
 income,
 credit_score,
 loan_amount,
 loan_duration,
 defaulted,
 CASE
 WHEN age < 30 THEN 'Young'
 WHEN age BETWEEN 30 AND 50 THEN 'Adult'
 ELSE 'Senior'
 END AS age_group,
 income / loan_amount AS income_to_loan_ratio,
 credit_score * loan_amount AS credit_score_loan_interaction
FROM credit_data;
```

- D.

```

CREATE OR REPLACE VIEW credit_data_enriched AS
SELECT ,
ASE
 WHEN age < 30 THEN 'Young'
 WHEN age BETWEEN 30 AND 50 THEN 'Adult'
 ELSE 'Senior'
AND AS age_group,
 income / loan_amount AS income_to_loan_ratio,
 credit_score * loan_amount AS credit_score_loan_interactio

```

• E. \_\_\_\_\_

**Answer: D,E**

Explanation:

Options D and E are correct. Option D creates a VIEW that dynamically calculates all three features without modifying the underlying table. A view is the correct and recommended usage. Option E creates a new table with all the features including the new engineered features. Option A creates a column and updates it, but this is inefficient compared to creating the feature directly in a single SELECT statement (Option E). B creates a temporary table but does not contain all three features. Option C, it only addresses the interaction feature, not age\_group or income to loan ratio.

#### NEW QUESTION # 225

You're tasked with building an image classification model on Snowflake to identify defective components on a manufacturing assembly line using images captured by high-resolution cameras. The images are stored in a Snowflake table named 'ASSEMBLY LINE IMAGES', with columns including 'image\_id' (INT), 'image\_data' (VARIANT containing binary image data), and 'timestamp' (TIMESTAMP NTZ). You have a pre-trained image classification model (TensorFlow/PyTorch) saved in Snowflake's internal stage. To improve inference speed and reduce data transfer overhead, which approach provides the MOST efficient way to classify these images using Snowpark Python and UDFs?

- A. Create a Python UDF that loads the entire table into memory, preprocesses the images, loads the pre-trained model, and performs classification for all images in a single execution.
- B. Create a Python UDF that takes a single 'image\_id' as input, retrieves the corresponding 'image\_data' from the table, preprocesses the image, loads the pre-trained model, performs classification, and returns the result. This UDF will be called for each image individually.
- C. Use Snowflake's external function feature to offload the image classification task to a serverless function hosted on AWS Lambda, passing the 'image\_id' to the function for processing.
- **D. Create a vectorized Python UDF that takes a batch of 'image\_id' values as input, retrieves the corresponding 'image\_data' from the 'ASSEMBLY LINE IMAGES' table using a JOIN, preprocesses the images in a vectorized manner, loads the pre-trained model once at the beginning, performs classification on the batch, and returns the results.**
- E. Create a Java UDF that loads the pre-trained model and preprocesses the images. Call this Java UDF from a Python UDF to perform the image classification. Since Java is faster than Python, this will optimize performance.

**Answer: D**

Explanation:

Option C offers the most efficient solution. Vectorized UDFs allow processing batches of data at once, significantly reducing overhead compared to processing each image individually (Option B). Loading the model once per batch avoids redundant model loading. Option A is highly inefficient as it attempts to load the entire table into memory. While Java can be faster in certain scenarios, the complexity of calling a Java UDF from a Python UDF (Option D) will likely introduce more overhead than benefits. External functions (Option E) introduce network latency and are generally less efficient than in-database processing, unless there's a specific need for external resources or specialized hardware that Snowflake doesn't offer.

#### NEW QUESTION # 226

You are developing a churn prediction model using Snowpark Python and Scikit-learn. After initial model training, you observe significant overfitting. Which of the following hyperparameter tuning strategies and code snippets, when implemented within a

Snowflake Python UDF, would be MOST effective to address overfitting in a Ridge Regression model and how can you implement a reproducible model with minimal code?

```
 Using 'GridSearchCV' with a wide range of 'alpha' values, without cross-validation, and then selecting the 'alpha' that gives the highest score on the training data. from sklearn.linear_model import Ridge; from sklearn.model_selection import GridSearchCV; param_grid = {'alpha': [0.001, 0.01, 0.1, 1, 10, 100]}; grid = GridSearchCV(Ridge(), param_grid, cv=None); grid.fit(X_train, y_train); best_alpha = grid.best_params_['alpha']

 Using 'RandomizedSearchCV' with a limited number of iterations and a fixed random state, along with cross-validation, and then selecting the 'alpha' that gives the highest average cross-validation score. from sklearn.linear_model import Ridge; from sklearn.model_selection import RandomizedSearchCV; from scipy.stats import loguniform; param_distributions = {'alpha': loguniform(1e-5, 100)}; rsearch = RandomizedSearchCV(Ridge(), param_distributions, n_iter=10, cv=3, random_state=42); rsearch.fit(X_train, y_train); best_alpha = rsearch.best_params_['alpha']

 Manually tuning the 'alpha' parameter by trial and error on the training data, without cross-validation or a structured search. from sklearn.linear_model import Ridge; alphas = [0.001, 0.01, 0.1, 1, 10, 100]; best_alpha = None; best_score = -float('inf'); for alpha in alphas: model = Ridge(alpha=alpha); model.fit(X_train, y_train); score = model.score(X_train, y_train); if score > best_score: best_score = score; best_alpha = alpha

 Using 'BayesianSearchCV' with Gaussian Processes and acquisition function optimization and a cross validation with n_jobs=-1. from sklearn.linear_model import Ridge; from skopt import BayesSearchCV; from skopt.space import Real; search_space = {'alpha': Real(1e-5, 100, prior='log-uniform')}; opt = BayesSearchCV(Ridge(), search_space, n_trials=5, cv=3, n_jobs=-1, random_state=42); opt.fit(X_train, y_train); best_alpha = opt.best_params_['alpha']

 Using 'HalvingGridSearchCV' with successive halving and resource allocation, without random state specification for complete reproducibility: from sklearn.experimental import enable_halving_search_cv # noqa; from sklearn.model_selection import HalvingGridSearchCV; from sklearn.linear_model import Ridge; param_grid = {'alpha': [0.001, 0.01, 0.1, 1, 10, 100]}; hsearch = HalvingGridSearchCV(Ridge(), param_grid, cv=3, resource='n_samples', max_resources=len(X_train), factor=3).fit(X_train, y_train); best_alpha = hsearch.best_params_['alpha']
```

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

Answer: A,D

Explanation:

Options B and D are correct because they employ techniques to mitigate overfitting. Option B uses 'RandomizedSearchCV' with cross-validation and a fixed 'random\_state', making the search reproducible and preventing overfitting by evaluating performance on multiple validation sets. Option D leverages 'BayesianSearchCV', which uses a probabilistic model to efficiently explore the hyperparameter space, also with cross-validation and a fixed random state making search reproducible. Both methods aim to find a balance between model complexity and generalization ability. Option A is incorrect because it does not use cross-validation, which is crucial for preventing overfitting. Option C is incorrect because manual tuning without a systematic search and cross-validation is prone to bias and overfitting. Finally, option E is incorrect because while using a modern algorithm, it lacks a random state, making it difficult to reproduce the outcome.

### NEW QUESTION # 227

.....

In accordance with the actual exam, we provide the latest DSA-C03 exam dumps for your practices. With the latest DSA-C03 test questions, you can have a good experience in practicing the test. Moreover, you have no need to worry about the price, we provide free updating for one year and half price for further partnerships, which is really a big sale in this field. After your payment, we will send the updated DSA-C03 Exam to you immediately and if you have any question about updating, please leave us a message.

**DSA-C03 Valid Exam Pattern:** <https://www.pass4suresvce.com/DSA-C03-pass4sure-vce-dumps.html>

Snowflake Updated DSA-C03 Testkings Thus your certification cost will be minimized, Professional experts Our professional experts are conversant about the practice materials, who are curious and careful specialists dedicated to better the DSA-C03 sure-pass learning materials: SnowPro Advanced: Data Scientist Certification Exam with diligence and outstanding knowledge all these years, Snowflake Updated DSA-C03 Testkings It will be your loss if you do not choose our study material.

Commodity specialists in the cavernous trading Updated DSA-C03 Testkings room are also caught off guard by the jobs report and are now gesturing wildly and barking out orders to buy oil and gasoline contracts DSA-C03 Valid Mock Test on the expectation that a resilient economy will drive up demand for fuel in the future.

## Latest Updated Updated DSA-C03 Testkings - Snowflake DSA-C03 Valid Exam Pattern: SnowPro Advanced: Data Scientist Certification Exam

The software is compatible with Windows so you can run it easily on DSA-C03 your computer, Thus your certification cost will be minimized, Professional experts Our professional experts are conversant about the practice materials, who are curious and careful

