

Actual4Dump



Pass Your Next Certification Exam Fast!

Everything you need to prepare, learn & pass your certification exam easily.

365 days free updates. First attempt guaranteed success.

| Choose the version that fits your needs | PDF Version | Desktop Test Engine | Online Test Engine |
|---|-------------|---------------------|--------------------|
| Latest and Up-to-Date exam dumps with real exam questions answers. | ✓ | ✓ | ✓ |
| Get 12-Months free updates without any extra charges. | ✓ | ✓ | ✓ |
| Experience same exam environment before appearing in the certification exam. | ✗ | ✓ | ✓ |
| 100% exam passing guarante in the first attempt. | ✓ | ✓ | ✓ |
| 20% discount on more than one license and 30% discount on 5+ license purchases. | ✗ | ✓ | ✓ |
| 100% secure purchase on SSL. | ✓ | ✓ | ✓ |
| Completely private purchase without sharing your personal info with anyone. | ✓ | ✓ | ✓ |

<http://www.actual4dump.com>

Superb Exam Dumps Materials lead you to get your certification easily - Actual4dump

Exam : **Databricks-Machine-Learning-Professional**

Title : Databricks Certified Machine Learning Professional

Vendor : Databricks

Version : DEMO

NO.1 A data scientist has developed a model and computed the RMSE of the model on the test set. They have assigned this value to the variable `rmse`. They now want to manually store the RMSE value with the MLflow run.

They write the following incomplete code block:

```
with mlflow.start_run(experiment_id=exp_id, run_name=run_name) as run:
    # Log rmse
    mlflow.____("rmse", rmse)
```

Which of the following lines of code can be used to fill in the blank so the code block can successfully complete the task?

- A. `log_artifact`
- B. `log_model`
- C. `log_metric`
- D. `log_param`
- E. There is no way to store values like this.

Answer: A

NO.2 Which of the following MLflow operations can be used to automatically calculate and log a Shapley feature importance plot?

- A. `mlflow.shap.log_explanation`
- B. None of these operations can accomplish the task.
- C. `mlflow.shap`
- D. `mlflow.log_figure`
- E. `client.log_artifact`

Answer: C

NO.3 A data scientist has developed a scikit-learn random forest model, but they have not yet logged model with MLflow. They want to obtain the input schema and the output schema of the model so they can document what type of data is expected as input. Which of the following MLflow operations can be used to perform this task?

- A. `mlflow.models.schema.infer_schema`
- B. `mlflow.models.signature.infer_signature`
- C. `mlflow.models.Model.get_input_schema`
- D. `mlflow.models.Model.signature`
- E. There is no way to obtain the input schema and the output schema of an unlogged model.

Answer: B

NO.4 A machine learning engineer and data scientist are working together to convert a batch deployment to an always-on streaming deployment. The machine learning engineer has expressed that rigorous data tests must be put in place as a part of their conversion to account for potential changes in data formats. Which of the following describes why these types of data type tests and checks are particularly important for streaming deployments?

- A. Because the streaming deployment is always on, all types of data must be handled without producing an error

- B. All of these statements
- C. Because the streaming deployment is always on, there is no practitioner to debug poor model performance
- D. Because the streaming deployment is always on, there is a need to confirm that the deployment can autoscale
- E. None of these statements

Answer: A

NO.5 Which deployment paradigm can centrally compute predictions for a single record with exceedingly fast results?

- A. Streaming
- B. Batch
- C. Edge/on-device
- D. None of these strategies will accomplish the task.
- E. Real-time

Answer: E

NO.6 A machine learning engineering team wants to build a continuous pipeline for data preparation of a machine learning application. The team would like the data to be fully processed and made ready for inference in a series of equal-sized batches. Which tool can be used to provide this type of continuous processing?

- A. Spark UDFs
- B. Structured Streaming
- C. MLflow
- D. Delta Lake

Answer: A

NO.7 A machine learning engineer wants to deploy a model for real-time serving using MLflow Model Serving. For the model, the machine learning engineer currently has one model version in each of the stages in the MLflow Model Registry. The engineer wants to know which model versions can be queried once Model Serving is enabled for the model. Which of the following lists all of the MLflow Model Registry stages whose model versions are automatically deployed with Model Serving?

- A. Staging, Production, Archived
- B. Production
- C. None, Staging, Production, Archived
- D. Staging, Production
- E. None, Staging, Production

Answer: D

NO.8 A data scientist has written a function to track the runs of their random forest model. The data scientist is changing the number of trees in the forest across each run. Which of the following MLflow operations is designed to log single values like the number of trees in a random forest?

- A. mlflow.log_artifact
- B. mlflow.log_model
- C. mlflow.log_metric
- D. mlflow.log_param
- E. There is no way to store values like this.

Answer: C

NO.9 A machine learning engineer is converting a Hyperopt-based hyperparameter tuning process from manual MLflow logging to MLflow Autologging. They are trying to determine how to manage nested Hyperopt runs with MLflow Autologging.

Which of the following approaches will create a single parent run for the process and a child run for each unique combination of hyperparameter values when using Hyperopt and MLflow Autologging?

- A. Starting a manual parent run before calling fmin
- B. Ensuring that a built-in model flavor is used for the model logging
- C. Starting a manual child run within the objective function
- D. There is no way to accomplish nested runs with MLflow Autologging and Hyperopt
- E. MLflow Autologging will automatically accomplish this task with Hyperopt

Answer: A

NO.10 A data scientist has created a Python function `compute_features` that returns a Spark DataFrame with the following schema:

```
customer_id STRING,  
spend DOUBLE,  
units INT,  
loyal INT,  
region STRING
```

The resulting DataFrame is assigned to the `features_df` variable. The data scientist wants to create a Feature Store table using `features_df`.

Which of the following code blocks can they use to create and populate the Feature Store table using the Feature Store Client `fs`?

A.

```
fs.create_table(  
    name="new_table",  
    primary_keys="customer_id",  
    df=features_df,  
    description="Customer features"  
)
```

B.

```
fs.create_table(  
    name="new_table",  
    primary_keys="customer_id",  
    description="Customer features"  
)
```

C. features_df.write.mode("fs").path("new_table")

D.

```
fs.create_table(  
    name="new_table",  
    primary_keys="customer_id",  
    function=compute_features,  
    description="Customer features"  
)
```

E. features_df.write.mode("feature").path("new_table")

Answer: D

NO.11 Which of the following is a benefit of logging a model signature with an MLflow model?

- A. The model will have a unique identifier in the MLflow experiment
- B. The schema of input data can be validated when serving models
- C. The model can be deployed using real-time serving tools
- D. The model will be secured by the user that developed it
- E. The schema of input data will be converted to match the signature

Answer: E

NO.12 Which statement describes streaming with Spark as a model deployment strategy?

- A. The inference of batch processed records as soon as a trigger is hit
- B. The inference of all types of records in real-time
- C. The inference of batch processed records as soon as a Spark job is run
- D. The inference of incrementally processed records as soon as trigger is hit
- E. The inference of incrementally processed records as soon as a Spark job is run

Answer: E

NO.13 A machine learning engineer has deployed a model recommender using MLflow Model Serving.

They now want to query the version of that model that is in the Production stage of the MLflow Model Registry. Which of the following model URIs can be used to query the described model version?

- A. `https://<databricks-instance>/model-serving/recommender/Production/invocations`
- B. The version number of the model version in Production is necessary to complete this task.
- C. `https://<databricks-instance>/model/recommender/stage-production/invocations`
- D. `https://<databricks-instance>/model-serving/recommender/stage-production/invocations`
- E. `https://<databricks-instance>/model/recommender/Production/invocations`

Answer: B

NO.14 Which of the following tools can assist in real-time deployments by packaging software with its own application, tools, and libraries?

- A. Cloud-based compute
- B. None of these tools
- C. REST APIs
- D. Containers
- E. Autoscaling clusters

Answer: A

NO.15 A machine learning engineer has registered a sklearn model in the MLflow Model Registry using the sklearn model flavor with UI model_uri. Which operation can be used to load the model as an sklearn object for batch deployment?

- A. `mlflow.spark.load_model(model_uri)`
- B. `mlflow.pyfunc.read_model(model_uri)`
- C. `mlflow.sklearn.read_model(model_uri)`
- D. `mlflow.pyfunc.load_model(model_uri)`
- E. `mlflow.sklearn.load_model(model_uri)`

Answer: D

NO.16 A data scientist set up a machine learning pipeline to automatically log a data visualization with each run. They now want to view the visualizations in Databricks. Which location in Databricks will show these data visualizations?

- A. The MLflow Model Registry Model page
- B. The Artifacts section of the MLflow Experiment page

- C. Logged data visualizations cannot be viewed in Databricks
- D. The Artifacts section of the MLflow Run page
- E. The Figures section of the MLflow Run page

Answer: E

NO.17 A data scientist has developed a scikit-learn model `sklearn_model` and they want to log the model using MLflow.

They write the following incomplete code block:

```
with mlflow.start_run(experiment_id=exp_id, run_name=run_name) as run:
    # Log model
    _____
```

Which lines of code can be used to fill in the blank so the code block can successfully complete the task?

- A. `mlflow.spark.track_model(sklearn_model, "model")`
- B. `mlflow.sklearn.log_model(sklearn_model, "model")`
- C. `mlflow.spark.log_model(sklearn_model, "model")`
- D. `mlflow.sklearn.load_model("model")`
- E. `mlflow.sklearn.track_model(sklearn_model, "model")`

Answer: A

NO.18 Which of the following describes the concept of MLflow Model flavors?

- A. A convention that deployment tools can use to wrap preprocessing logic into a Model
- B. A convention that MLflow Model Registry can use to version models
- C. A convention that MLflow Experiments can use to organize their Runs by project
- D. A convention that deployment tools can use to understand the model
- E. A convention that MLflow Model Registry can use to organize its Models by project

Answer: C

NO.19 In a continuous integration, continuous deployment (CI/CD) process for machine learning pipelines, which of the following events commonly triggers the execution of automated testing?

- A. The launch of a new cost-efficient SQL endpoint
- B. CI/CD pipelines are not needed for machine learning pipelines
- C. The arrival of a new feature table in the Feature Store
- D. The launch of a new cost-efficient job cluster
- E. The arrival of a new model version in the MLflow Model Registry

Answer: D