

Dumps4u

Vendor: Microsoft

Exam Code: 70-475

Exam Name: Designing and Implementing Big Data Analytics Solutions

Version: DEMO

More Information: <https://www.dumps4u.com/70-475-dumps/>



24X7 X 365
Customer Support



Instant Download
Right After Purchase



SECURE SITE
100% secure site

Product Version: 6

Case Study: 1

Mix Questions

Question: 1

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

Your company has multiple databases that contain millions of sales transactions.
You plan to implement a data mining solution to identify purchasing fraud.

<https://www.dumps4u.com/>

You need to design a solution that mines 10 terabytes (TB) of sales dat

a. The solution must meet the following requirements:

- Run the analysis to identify fraud once per week.
- Continue to receive new sales transactions while the analysis runs.
- Be able to stop computing services when the analysis is NOT running.

Solution: You create a Cloudera Hadoop cluster on Microsoft Azure virtual machines.

Does this meet the goal?

A. Yes

B. No

Answer: A

Question: 2

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

Your company has multiple databases that contain millions of sales transactions.

You plan to implement a data mining solution to identify purchasing fraud.

You need to design a solution that mines 10 terabytes (TB) of sales dat

a. The solution must meet the following requirements:

Run the analysis to identify fraud once per week.

Continue to receive new sales transactions while the analysis runs.

Be able to stop computing services when the analysis is NOT running.

Solution: You create a Microsoft Azure Data Lake job.

Does this meet the goal?

A. Yes

B. No

Answer: A

Question: 3

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

Your company has multiple databases that contain millions of sales transactions.

You plan to implement a data mining solution to identify purchasing fraud.

You need to design a solution that mines 10 terabytes (TB) of sales dat

a. The solution must meet the following requirements:

Run the analysis to identify fraud once per week.

Continue to receive new sales transactions while the analysis runs.

Be able to stop computing services when the analysis is NOT running.

Solution: You create a Microsoft Azure HDInsight cluster.

Does this meet the goal?

A. Yes

B. No

Answer: B

Question: 4

HOTSPOT

You are designing a solution based on the lambda architecture.

The solution has the following layers;

Batch

Speed

Serving

You are planning the data ingestion process and the query execution.

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer Area

The data ingestion process must only communicate with the batch layer:

The query execution must communicate with both the serving layer and the speed layer:

You can use Kafka to execute the queries:

Answer:

Answer Area

The data ingestion process must only communicate with the batch layer:

The query execution must communicate with both the serving layer and the speed layer:

You can use Kafka to execute the queries:

Question: 5

DRAG DROP

You have a web app that accepts user input, and then uses a Microsoft Azure Machine Learning model to predict a characteristic of the user.

You need to perform the following operations:

Track the number of web app users from month to month.

Track the number of successful predictions made during the last minute.

Create a dashboard showcasing the analytics for the predictions and the web app usage.

Which lambda layer should you query for each operation? To answer, drag the appropriate layers to the correct operations. Each layer may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Layers	Answer Area
Batch	Track the number of successful predictions made during the last minute: Layers
Speed	Track the number of web app users from month to month: Layers
Serving	Create a dashboard showcasing the analytics for the predictions and the web app usage: Layers

Answer:

Track the number of successful predictions made during the last minute: Serving

Track the number of web app users from month to month: Batch

Create a dashboard showcasing the analytics for the predictions and the web app usage: Speed

THANKS FOR TRYING OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual [70-475](#)
Exam Questions With Answers.

Dumps4u also provides you extra [20% discount](#) on the purchase
of 70-475 exam. Use Coupon code: "[dumps4u](#)"

More Information: <https://www.dumps4u.com/70-475-dumps/>

Money Back Guarantee



<https://www.dumps4u.com/>