



70-475^{Q&As}

Designing and Implementing Big Data Analytics Solutions

Pass Microsoft 70-475 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.lead4pass.com/70-475.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

Your company has several thousand sensors deployed.

You have a Microsoft Azure Stream Analytics job that receives two data streams Input1 and Input2 from an Azure event hub. The data streams are portioned by using a column named SensorName. Each sensor is identified by a field named

SensorID.

You discover that Input2 is empty occasionally and the data from Input1 is ignored during the processing of the Stream Analytics job.

You need to ensure that the Stream Analytics job always processes the data from Input1.

How should you modify the query? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
SELECT I1.SensorID, I1.EntryTime, I2.ExitTime, [DATEDIFF  
(minute, I1.EntryTime, I2.ExitTime)] as Duration  
FROM Input I1 TIMESTAMP BY EntryTime PARTITION BY SensorName
```

	▼
JOIN	
LEFT OUTER JOIN	
RIGHT OUTER JOIN	

```
Input I2 TIMESTAMP BY ExitTime PARTITION BY SensorName
```

```
JOIN table1 R  
ON R.SensorID = I1.SensorID  
AND [DATEDIFF(minute, I1, I2) BETWEEN 0 AND 10]  
WHERE R.IsValid = 1
```

	▼
ON I1.SensorId=I2.SensorID	
ON I1.SensorName=I2.SensorName	
ON I1.SensorId=I2.SensorId OR I1.SensorName=I2.SensorName	

Correct Answer:



Answer Area

```
SELECT I1.SensorID, I1.EntryTime, I2.ExitTime, [DATEDIFF
(minute, I1.EntryTime, I2.ExitTime)] as Duration
FROM Input I1 TIMESTAMP BY EntryTime PARTITION BY SensorName
```

JOIN

LEFT OUTER JOIN

RIGHT OUTER JOIN

```
Input I2 TIMESTAMP BY ExitTime PARTITION BY SensorName
```

```
JOIN table1 R
ON R.SensorID = I1.SensorID
AND [DATEDIFF(minute, I1, I2) BETWEEN 0 AND 10]
WHERE R.IsValid = 1
```

ON I1.SensorId=I2.SensorID

ON I1.SensorName=I2.SensorName

ON I1.SensorId=I2.SensorId OR I1.SensorName=I2.SensorName

Box 1: LEFT OUTER JOIN LEFT OUTER JOIN specifies that all rows from the left table not meeting the join condition are included in the result set, and output columns from the other table are set to NULL in addition to all rows returned by the inner join. Box 2: ON I1.SensorID=I2.SensorID References: <https://docs.microsoft.com/en-us/stream-analytics-query/join-azure-stream-analytics>

QUESTION 2

A company named Fabrikam, Inc. has a web app. Millions of users visit the app daily.

Fabrikam performs a daily analysis of the previous day's logs by scheduling the following Hive query.

```
CREATE EXTERNAL TABLE IF NOT EXISTS UserActivity (...) PARTITIONED BY (LogDate string) LOCATION 'wasb:///logs';
MSCK REPAIR TABLE UserActivity;
Select ... From UserActivity where LogDate = "{date}";
```

You need to recommend a solution to gather the log collections from the web app. What should you recommend?

- A. Generate a single directory that has a single file for each day.
- B. Generate a single directory that contains multiple files for each day. Name the file by using the syntax of {date}_{randomsuffix}.txt.
- C. Generate a directory each day that has a single file.
- D. Generate a directory that is named by using the syntax of "LogDate={date}" and generate a set of files for that day.

Correct Answer: D



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

the 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.

You have an Apache Spark system that contains 5 TB of data.

You need to write queries that analyze the data in the system. The queries must meet the following requirements:

Use static data typing.

Execute queries as quickly as possible.

Have access to the latest language features.

Solution: You write the queries by using Scala.

A. Yes

B. No

Correct Answer: A

QUESTION 4

You have a Microsoft Azure HDInsight cluster for analytics workloads. You have a C# application on a local computer. You plan to use Azure Data Factory to run the C# application in Azure.

You need to create a data factory that runs C# application by using HDInsight.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order. NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Select and Place:

Actions

- Derive the C# class from the IDotNetActivity interface.
- Zip the build files and upload the ZIP file to an Azure storage account.
- Implement the Execute method of the IDotNetActivity interface.
- Build the C# application in Microsoft Visual Studio.
- Create a pipeline that has a DotNetActivity activity and specify the path to the build files in the Azure storage account.

Answer Area





Correct Answer:

Actions	Answer Area
	Derive the C# class from the IDotNetActivity interface.
	Implement the Execute method of the IDotNetActivity interface.
	Build the C# application in Microsoft Visual Studio.
	Zip the build files and upload the ZIP file to an Azure storage account.
	Create a pipeline that has a DotNetActivity activity and specify the path to the build files in the Azure storage account.

Step 1: Derive the C# class from the IDotNetActivity interface

Step 2: Implement the Execute method of the IDotNetActivity interface.

Implement (Add) the Execute method of the IDotNetActivity interface to the MyDotNetActivity class

Step 3: Build the C# application in Microsoft Visual Studio

Compile the project. Click Build from the menu and click Build Solution.

Step 4: Zip the build files and upload the ZIP file to an Azure storage account.

Create a zip file MyDotNetActivity.zip that contains all the binaries in the project folder. Include the MyDotNetActivity.pdb file so that you get additional details such as line number in the source code that caused the issue if there was a failure.

Step 5: Create a pipeline..

You have created a custom activity and uploaded the zip file with binaries to a blob container in a general-purpose Azure Storage Account. Now you create an Azure data factory with a pipeline that uses the custom activity.

References:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/data-factory/v1/data-factory-use-custom-activities.md>

QUESTION 5

You are designing a partitioning scheme for ingesting real-time data by using Kafka. Kafka and Apache Storm will be integrated.

You plan to use four event processing servers that each run as a Kafka consumer. Each server will have a two quad-core processor.

You need to identify the minimum number of partitions required to ensure that the load is distributed evenly.

How many should you identify?

A. 1

B. 4



C. 16

D. 32

Correct Answer: C

[Latest 70-475 Dumps](#)

[70-475 PDF Dumps](#)

[70-475 Study Guide](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate

More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)


We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.lead4pass.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 <p>One Year Free Update Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p>Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p>Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © lead4pass, All Rights Reserved.