



70-776^{Q&As}

Perform Big Data Engineering on Microsoft Cloud Services

Pass Microsoft 70-776 Exam with 100% Guarantee

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

<https://www.passapply.com/70-776.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

You ingest data into a Microsoft Azure event hub.

You need to export the data from the event hub to Azure Storage and to prepare the data for batch processing tasks in Azure Data Lake Analytics.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Run the Avro extractor from a U-SQL script.
- B. Create an Azure Storage account.
- C. Add a shared access policy.
- D. Enable Event Hubs Archive.
- E. Run the CSV extractor from a U-SQL script.

Correct Answer: BD

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.

You have a table named Table1 that contains 3 billion rows. Table1 contains data from the last 36 months.

At the end of every month, the oldest month of data is removed based on a column named DateTime.

You need to minimize how long it takes to remove the oldest month of data.

Solution: You implement a columnstore index on the DateTime column.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A

QUESTION 3



You are using Cognitive capabilities in U-SQL to analyze images that contain different types of objects.

You need to identify which objects might be people.

Which two reference assemblies should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. ExtPython
- B. ImageCommon
- C. ImageTagging
- D. ExtR
- E. FaceSdk

Correct Answer: BC

QUESTION 4

You have a Microsoft Azure SQL data warehouse.

You need to configure Data Warehouse Units (DWUs) to ensure that you have six compute nodes. The solution must minimize costs.

Which value should set for the DWUs?

- A. DW200
- B. DW400
- C. DW600
- D. DW1000

Correct Answer: C

QUESTION 5

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.

You have a table named Table1 that contains 3 billion rows. Table1 contains data from the last 36 months.

At the end of every month, the oldest month of data is removed based on a column named DateTime.



You need to minimize how long it takes to remove the oldest month of data.

Solution: You specify DateTime as the hash distribution column.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

QUESTION 6

DRAG DROP

You have an Apache Hive database in a Microsoft Azure HDInsight cluster.

You create an Azure Data Factory named DF1.

You need to transform the data in the Hive database and to output the data to Azure Blob storage.


Which three cmdlets should you run in sequence? To answer, move the appropriate cmdlets from the list of cmdlets to the answer area and arrange them in the correct order.

Select and Place:

Cmdlets

- Set-AzureRmDataFactoryGateway
- New-AzureRmDataFactoryLinkedService
- New-AzureRmDataFactoryDataset
- New-AzureRmDataFactoryGateway
- Set-AzureRmDataFactory
- New-AzureRmDataFactoryHub
- New-AzureRmDataFactoryPipeline

Answer Area



Correct Answer:



Cmdlets

Answer Area



References: <https://docs.microsoft.com/en-us/powershell/module/azurermdatadf/new-azurermdatadfdatapipeline?view=azurermps-4.4.0> <https://github.com/aelij/azure-content/blob/master/articles/data-factory/data-factory-build-your-first-pipeline-using-powershell.md>

QUESTION 7

HOTSPOT

You use Microsoft Visual Studio to develop custom solutions for customers who use Microsoft Azure Data Lake Analytics.

You install the Data Lake Tools for Visual Studio.

You need to identify which tasks can be performed from Visual Studio and which tasks can be performed from the Azure portal.

What should you identify for each task? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

Create a U-SQL project:

Azure portal only
Visual Studio only
Both Visual Studio and Azure portal

Create an Azure Data Lake Analytics account:

Azure portal only
Visual Studio only
Both Visual Studio and Azure portal

Upload data files to the default Azure Data Lake Store:

Azure portal only
Visual Studio only
Both Visual Studio and Azure portal

Correct Answer:

Answer Area

Create a U-SQL project:

Azure portal only
Visual Studio only
Both Visual Studio and Azure portal

Create an Azure Data Lake Analytics account:

Azure portal only
Visual Studio only
Both Visual Studio and Azure portal

Upload data files to the default Azure Data Lake Store:

Azure portal only
Visual Studio only
Both Visual Studio and Azure portal

QUESTION 8**HOTSPOT**

You plan to implement a Microsoft Azure Stream Analytics job to track the data from IoT devices. You will have the following two jobs:

- Job1 will contain a query that has one non-partitioned step.
- Job2 will contain a query that has two steps. One of the steps is partitioned.



What is the maximum number of streaming units that will be consumed per job? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Job1:

1
3
6
12
18
24
30
32
42
48

Job2:

1
3
6
12
18
24
30
32
42
48

Correct Answer:



Answer Area

Job1: ▼

1
3
6
12
18
24
30
32
42
48

Job2: ▼

1
3
6
12
18
24
30
32
42
48

www.PassApply.com

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-scale-jobs> <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-streaming-unit-consumption>

QUESTION 9

You have a Microsoft Azure Data Lake Analytics service and an Azure Data Lake Store.

You need to use Python to submit a U-SQL job.

Which Python module should you install?

- A. azure-mgmt-datalake-store
- B. azure-mgmt-datalake-analytics
- C. azure-datalake-store
- D. azure-mgmt-resource

Correct Answer: B

QUESTION 10

You have a Microsoft Azure SQL data warehouse that has a fact table named FactOrder. FactOrder contains three columns named CustomerId, OrderId, and OrderDateKey. FactOrder is hash distributed on CustomerId. OrderId is the unique



identifier for FactOrder.

FactOrder contains 3 million rows.

Orders are distributed evenly among different customers from a table named dimCustomers that contains 2 million rows.

You often run queries that join FactOrder and dimCustomers by selecting and grouping by the OrderDateKey column.

You add 7 million rows to FactOrder. Most of the new records have a more recent OrderDateKey value than the previous records.

You need to reduce the execution time of queries that group on OrderDateKey and that join dimCustomers and FactOrder.

What should you do?

- A. Change the distribution for the FactOrder table to round robin.
- B. Update the statistics for the OrderDateKey column.
- C. Change the distribution for the FactOrder table to be based on OrderId.
- D. Change the distribution for the dimCustomers table to OrderDateKey.

Correct Answer: B

QUESTION 11

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.

You have a table named Table1 that contains 3 billion rows. Table1 contains data from the last 36 months.

At the end of every month, the oldest month of data is removed based on a column named DateTime.

You need to minimize how long it takes to remove the oldest month of data.

Solution: You implement round robin for table distribution. Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

QUESTION 12



You plan to deploy a Microsoft Azure virtual machine that will host a data warehouse. The data warehouse will contain a 10-TB database.

You need to provide the fastest read and write times for the database.

Which disk configuration should you use?

- A. storage pools with mirrored disks
- B. RAID 5 volumes
- C. spanned volumes
- D. striped volumes
- E. storage pools with striped disks

Correct Answer: E

[Latest 70-776 Dumps](#)

[70-776 Practice Test](#)

[70-776 Exam Questions](#)



VCE & PDF

PassApply.com

<https://www.passapply.com/70-776.html>

2021 Latest passapply 70-776 PDF and VCE dumps Download

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.passapply.com/allproducts>

Need Help

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

To update a previously submitted ticket:



 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.	 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.	 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.
---	---	--

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

All trademarks are the property of their respective owners.

Copyright © passapply, All Rights Reserved.