# Microsoft

## 70-776 Exam

**Microsoft Perform Big Data Engineering on Microsoft Cloud Services (beta) Exam**

# Version: 10.0

## 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.

You are monitoring user queries to a Microsoft Azure SQL data warehouse that has six compute nodes.

You discover that compute node utilization is uneven. The rows_processed column from sys.dm_pdw_workers shows a significant variation in the number of rows being moved among the distributions for the same table for the same query.

You need to ensure that the load is distributed evenly across the compute nodes.

Solution: You add a clustered columnstore index.

Does this meet the goal?

A. Yes
B. No

**Answer: B**

## 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 are monitoring user queries to a Microsoft Azure SQL data warehouse that has six compute nodes.

You discover that compute node utilization is uneven. The rows_processed column from sys.dm_pdw_workers shows a significant variation in the number of rows being moved among the distributions for the same table for the same query.

You need to ensure that the load is distributed evenly across the compute nodes.

Solution: You add a nonclustered columnstore index.

Does this meet the goal?

A. Yes
B. No

**Answer: B**

## 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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are monitoring user queries to a Microsoft Azure SQL data warehouse that has six compute nodes.

You discover that compute node utilization is uneven. The rows_processed column from sys.dm_pdw_workers shows a significant variation in the number of rows being moved among the distributions for the same table for the same query.

You need to ensure that the load is distributed evenly across the compute nodes.

Solution: You change the table to use a column that is not skewed for hash distribution.

Does this meet the goal?

A. Yes
B. No

**Answer: A**

## Question: 4

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

**Answer: B**

## 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 implement a columnstore index on the DateTime column.

Does this meet the goal?

A. Yes
B. No

**Answer: A**

## Question: 6

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

**Answer: B**

## Question: 7

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 are troubleshooting a slice in Microsoft Azure Data Factory for a dataset that has been in a waiting state for the last three days. The dataset should have been ready two days ago.

The dataset is being produced outside the scope of Azure Data Factory. The dataset is defined by using the following JSON code.

```json
{
"name": "CustomerTable",
"properties": {
"type": "AzureBlob",
"linkedServiceName": "MyLinkedService",
"typeProperties": {
"folderPath": "MyContainer/MySubFolder/",
"format": {
"type": "TextFormat",
"columnDelimiter": ",",
"rowDelimiter": ";"
}
},
"external": false,
"availability": {
"frequency": "Hour",
"interval": 1
},
"policy": {
}
}
}
}
```

You need to modify the JSON code to ensure that the dataset is marked as ready whenever there is data in the data store.
Solution: You change the external attribute to true.
Does this meet the goal?

A. Yes
B. No

---

**Answer: A**

References:
https://docs.microsoft.com/en-us/azure/data-factory/v1/data-factory-create-datasets

## Question: 8

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 are troubleshooting a slice in Microsoft Azure Data Factory for a dataset that has been in a waiting state for the last three days. The dataset should have been ready two days ago.

The dataset is being produced outside the scope of Azure Data Factory. The dataset is defined by using the following JSON code.

```
{
"name": "CustomerTable",
"properties": {
"type": "AzureBlob",
"linkedServiceName": "MyLinkedService",
"typeProperties": {
"folderPath": "MyContainer/MySubFolder/",
"format": {
"type": "TextFormat",
"columnDelimiter": ",",
"rowDelimiter": ";"
}
},
"external": false,
"availability": {
"frequency": "Hour",
"interval": 1
},
"policy": {
}
}
}
}
```

You need to modify the JSON code to ensure that the dataset is marked as ready whenever there is data in the data store.
Solution: You change the interval to 24.
Does this meet the goal?

A. Yes
B. No

**Answer: B**

References:
https://docs.microsoft.com/en-us/azure/data-factory/v1/data-factory-create-datasets

## Question: 9

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 are troubleshooting a slice in Microsoft Azure Data Factory for a dataset that has been in a waiting state for the last three days. The dataset should have been ready two days ago.

The dataset is being produced outside the scope of Azure Data Factory. The dataset is defined by using the following JSON code.

```
{
"name": "CustomerTable",
"properties": {
"type": "AzureBlob",
"linkedServiceName": "MyLinkedService",
"typeProperties": {
"folderPath": "MyContainer/MySubFolder/",
"format": {
"type": "TextFormat",
"columnDelimiter": ",",
"rowDelimiter": ";"
}
},
"external": false,
"availability": {
"frequency": "Hour",
"interval": 1
},
"policy": {
}
}
}
}
```

You need to modify the JSON code to ensure that the dataset is marked as ready whenever there is data in the data store.
Solution: You add a structure property to the dataset.
Does this meet the goal?

A. Yes
B. No

**Answer: B**

References:
https://docs.microsoft.com/en-us/azure/data-factory/v1/data-factory-create-datasets

**Question: 10**

Note: This question is part of a series of questions that present the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.
Start of repeated scenario
You are migrating an existing on-premises data warehouse named LocalDW to Microsoft Azure. You will use an Azure SQL data warehouse named AzureDW for data storage and an Azure Data Factory named AzureDF for extract, transformation, and load (ETL) functions.
For each table in LocalDW, you create a table in AzureDW.
On the on-premises network, you have a Data Management Gateway.
Some source data is stored in Azure Blob storage. Some source data is stored on an on-premises Microsoft SQL Server instance. The instance has a table named Table1.
After data is processed by using AzureDF, the data must be archived and accessible forever. The archived data must meet a Service Level Agreement (SLA) for availability of 99 percent. If an Azure region fails, the archived data must be available for reading always. The storage solution for the archived data must minimize costs.
End of repeated scenario.
You need to configure an activity to move data from blob storage to AzureDW.
What should you create?

A. a pipeline
B. a linked service
C. an automation runbook
D. a dataset

**Answer: A**

References:
https://docs.microsoft.com/en-us/azure/data-factory/v1/data-factory-azure-blob-connector

**Thank You for trying 70-776 PDF Demo**

To Buy Latest 70-776 Full Version Download visit link below

https://www.certkillers.net/Exam/70-776

# Start Your 70-776 Preparation

*[Limited Time Offer]* Use Coupon "CKNET" for Further  discount on your purchase. Test your 70-776 preparation with actual exam questions.