

## 70-776 Dumps

### Perform Big Data Engineering on Microsoft Cloud Services (beta)

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



#### NEW 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**

#### NEW 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**

#### NEW 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.

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**

#### NEW 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 implement a columnstore index on the DateTime column. Does this meet the goal?

- A. Yes
- B. No

**Answer: A**

#### NEW 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 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 external attribute to true. Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

References:

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

**NEW 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 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

**Explanation:**

References:

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

**NEW 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.

```
{
  "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**

**Explanation:**

References:

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

**NEW QUESTION 8**

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**

**Explanation:**

References:

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

**NEW QUESTION 9**

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.

End of repeated scenario.

You need to configure Azure Data Factory to connect to the on-premises SQL Server instance. What should you do first?

- A. Deploy an Azure virtual network gateway.
- B. Create a dataset in Azure Data Factory.
- C. From Azure Data Factory, define a data gateway.
- D. Deploy an Azure local network gateway.

**Answer: C**

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/data-factory/v1/data-factory-move-data-between-onprem- and-cloud>

**NEW 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 define the schema of Table1 in AzureDF. What should you create?

- A. a gateway
- B. a linked service
- C. a dataset
- D. a pipeline

**Answer: C**

**NEW QUESTION 10**

HOTSPOT

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.

How should you configure the storage to archive the source data? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

Storage tier:

- Blob Storage Cool
- Blob Storage Hot
- General Purpose

Storage account type:

- Geo-Redundant Storage (GRS)
- Locally Redundant Storage (LRS)
- Read-Access Geo-Redundant Storage (RA-GRS)
- Zone-Redundant Storage (ZRS)

**Answer:**

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

**NEW QUESTION 12**

DRAG DROP

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.

Which three actions should you perform in sequence to migrate the on-premises data warehouse to Azure SQL Data Warehouse? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Create external tables and insert data into the AzureDW tables.	
Upload the files by using AzCopy.	
Upload files by using FTP.	⬆
Execute the CREATE TABLE AS SELECT statement.	⬆
Export the data to text files by using bcp.	
Import the files running the BULK INSERT command.	
Export data by using sqlcmd.	

**Answer:**

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-load-from-sql-server-with-polybase>

**NEW QUESTION 14**

DRAG DROP

You plan to develop a solution for real-time sentiment analysis of Twitter data.

You need to create a Microsoft Azure Stream Analytics job query to count the number of tweets during a period.

Which Window function should you use for each requirement? To answer, drag the appropriate functions to the correct requirements. Each function 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.

Windows Functions	Answer Area
HoppingWindow	Count the number of tweets every 10 seconds during the last 90 seconds: Window function
SlidingWindow	Count the number of tweets during the last 30 seconds: Window function
TumblingWindow	

**Answer:**

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-window-functions>

**NEW QUESTION 15**

DRAG DROP

You use Microsoft Azure Stream Analytics to analyze data from an Azure event hub in real time and send the output to a table named Table1 in an Azure SQL database. Table1 has three columns named Date, EventID, and User.

You need to prevent duplicate data from being stored in the database.

How should you complete the statement? To answer, drag the appropriate values to the correct targets. Each value 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.

Values	Answer Area
CHECK	ALTER TABLE Table1
CONSTRAINT	ADD [Value] Var1 [Value] (Date, EventID, User);
FOREIGN KEY	
PRIMARY KEY	
UNIQUE	

**Answer:**

**Explanation:**

Values	Answer Area
CHECK	ALTER TABLE Table1
CONSTRAINT	ADD CONSTRAINT Var1 UNIQUE (Date, EventID, User);
FOREIGN KEY	
PRIMARY KEY	
UNIQUE	

**NEW QUESTION 16**

You are building a Microsoft Azure Stream Analytics job definition that includes inputs, queries, and outputs. You need to create a job that automatically provides the highest level of parallelism to the compute instances. What should you do?

- A. Configure event hubs and blobs to use the PartitionKey field as the partition ID.
- B. Set the partition key for the inputs, queries, and outputs to use the same partition folder
- C. Configure the queries to use uniform partition keys.
- D. Set the partition key for the inputs, queries, and outputs to use the same partition folder
- E. Configure the queries to use different partition keys.
- F. Define the number of input partitions to equal the number of output partitions.

**Answer:** A

**Explanation:**

References:  
<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-parallelization>

**NEW QUESTION 21**

You are developing an application that uses Microsoft Azure Stream Analytics. You have data structures that are defined dynamically. You want to enable consistency between the logical methods used by stream processing and batch processing. You need to ensure that the data can be integrated by using consistent data points. What should you use to process the data?

- A. a vectorized Microsoft SQL Server Database Engine
- B. directed acyclic graph (DAG)
- C. Apache Spark queries that use updateStateByKey operators
- D. Apache Spark queries that use mapWithState operators

**Answer:** D

**NEW QUESTION 26**

DRAG DROP

You have a Microsoft Azure Stream Analytics solution that captures website visits and user interactions on the website. You have the sample input data described in the following table.

username	feature	EventType	EventTime
User1@contoso.com	Shopping cart	Start	2017-01-01T00:00:01.0000000Z
User1@contoso.com	Shopping cart	End	2017-01-01T00:00:01.0000000Z

You have the sample output described in the following table.

username	feature	duration in sec
User1@contoso.com	Shopping cart	7

How should you complete the script? To answer, drag the appropriate values to the correct targets. Each value 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.

Values	Answer Area
DATEADD	SELECT
DURATION	username,
LAG	feature,
DATEDIFF	Value (second,
FIRST(EventTime)	Value OVER (
LAST(EventTime)	PARTITION BY username,
	feature LIMIT Value (hour, 1) WHEN EventType = 'start'),
	EventTime)
	as 'duration in sec'
	FROM input TIMESTAMP BY Time
	WHERE EventType = 'end'

**Answer:**

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-stream-analytics-query-patterns>

**NEW QUESTION 31**

You have a Microsoft Azure SQL data warehouse that contains information about community events. An Azure Data Factory job writes an updated CSV file in Azure Blob storage to Community/{date}/events.csv daily.

You plan to consume a Twitter feed by using Azure Stream Analytics and to correlate the feed to the community events.

You plan to use Stream Analytics to retrieve the latest community events data and to correlate the data to the Twitter feed data.

You need to ensure that when updates to the community events data is written to the CSV files, the Stream Analytics job can access the latest community events data.

What should you configure?

- A. an output that uses a blob storage sink and has a path pattern of Community/{date}
- B. an output that uses an event hub sink and the CSV event serialization format
- C. an input that uses a reference data source and has a path pattern of Community/{date}/events.csv
- D. an input that uses a reference data source and has a path pattern of Community/{date}

**Answer: C**

**NEW QUESTION 34**

You plan to add a file from Microsoft Azure Data Lake Store to Azure Data Catalog. You run the Data Catalog tool and select Data Lake Store as the data source. Which information should you enter in the Store Account field to connect to the Data Lake Store?

- A. an email alias
- B. a server name
- C. a URL
- D. a subscription ID

**Answer: C**

**NEW QUESTION 36**

You have a Microsoft Azure Data Lake Analytics service.

You have a CSV file that contains employee salaries.

You need to write a U-SQL query to load the file and to extract all the employees who earn salaries that are greater than \$100,000. You must encapsulate the data for reuse.

What should you use?

- A. a table-valued function
- B. a view
- C. the extract command
- D. the output command

**Answer: A**

**Explanation:**

References:

<https://docs.microsoft.com/en-au/azure/data-lake-analytics/data-lake-analytics-u-sql-catalog>

**NEW QUESTION 38**

You have a Microsoft Azure Data Lake Analytics service. You plan to configure diagnostic logging.

You need to use Microsoft Operations Management Suite (OMS) to monitor the IP addresses that are used to access the Data Lake Store.

What should you do?

- A. Stream the request logs to an event hub.
- B. Send the audit logs to Log Analytics.
- C. Send the request logs to Log Analytics.
- D. Stream the audit logs to an event hub.

**Answer: B**

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/data-lake-analytics/data-lake-analytics-diagnostic-logs> <https://docs.microsoft.com/en-us/azure/security/azure-log-audit>

**NEW QUESTION 42**

You plan to use Microsoft Azure Data factory to copy data daily from an Azure SQL data warehouse to an Azure Data Lake Store.

You need to define a linked service for the Data Lake Store. The solution must prevent the access token from expiring.

Which type of authentication should you use?

- A. OAuth
- B. service-to-service

- C. Basic
- D. service principal

**Answer:** D

**Explanation:**

References:

<https://docs.microsoft.com/en-gb/azure/data-factory/v1/data-factory-azure-datalake-connector#azure-data-lake-store-linked-service-properties>

**NEW QUESTION 45**

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.

**Answer:** BD

**NEW QUESTION 48**

You have a Microsoft Azure Data Lake Analytics service.

You need to write a U-SQL query to extract from a CSV file all the users who live in Boston, and then to save the results in a new CSV file.

Which U-SQL script should you use?

**A**

```
@users =  
EXTRACT Name      string,  
         Age       int,  
         City      string,  
         State     string  
FROM "/users.csv"  
USING Extractors.Csv();  
  
@BostonUsers =  
SELECT Name, Age, City, State  
FROM @users  
WHERE City == "Boston";  
  
OUTPUT @BostonUsers  
      TO "/output/BostonUsers.csv"  
      USING Outputters.Csv();
```

**B**

```
@users =  
EXTRACT UserName  string,  
         Age       int,  
         City      string,  
         State     string  
FROM '/users.csv'  
USING Outputters.Csv();  
  
@BostonUsers =  
SELECT UserName, Age, City, State  
FROM @users  
WHERE City == "Boston";  
  
OUTPUT @BostonUsers  
      TO '/output/BostonUsers.csv'  
      USING Extractors.Csv();
```

C

```
@users =  
EXTRACT Name      string,  
        Age       int,  
        City      string,  
        State     string  
FROM '/users.csv'  
USING Extractors.Csv();  
  
@BostonUsers =  
SELECT UserName, Age, City, State  
FROM @users  
WHERE City == "Boston";  
  
OUTPUT @BostonUsers  
      TO '/output/BostonUsers.csv'  
      USING Outputters.Csv();
```

D

```
@users =  
EXTRACT UserName  string,  
        Age       int,  
        City      string,  
        State     string  
From "/users.csv"  
Using Extractors.Csv();  
  
@BostonUsers =  
SELECT UserName, Age, City, State  
From @users  
Where City == "Boston";  
  
OUTPUT @BostonUsers  
      TO "/output/BostonUsers.csv"  
      Using Outputters.Csv();
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

**Answer:** A

**NEW QUESTION 49**

You have a file in a Microsoft Azure Data Lake Store that contains sales data. The file contains sales amounts by salesperson, by city, and by state. You need to use U-SQL to calculate the percentage of sales that each city has for its respective state. Which code should you use?

**A**

```
@result=
SELECT
    City, State,
    SUM(SalesAmount)
        OVER( PARTITION BY City ) / SUM(SalesAmount)
        OVER( PARTITION BY State )
    AS CitySalesPercent
FROM @Sales;
```

**B**

```
@result=
SELECT City, SUM(SalesAmount)
AS CitySalesPercent
FROM @Sales;
GROUP BY City;
```

**C**

```
@result=
SELECT
    Salesperson, City, State,
    SUM(SalesAmount)
        OVER( PARTITION BY City ) / SUM(SalesAmount)
        OVER()
    AS CitySalesPercent
FROM @Sales;
```

**D**

```
@result=
SELECT
    City, State,
    SUM(SalesAmount)
        OVER( ) / SUM(SalesAmount)
        OVER( )
    AS CitySalesPercent
FROM @Sales;
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

**Answer:** A

**NEW QUESTION 54**

You are developing an application by using the Microsoft .NET SDK. The application will access data from a Microsoft Azure Data Lake folder. You plan to authenticate the application by using service-to-service authentication. You need to ensure that the application can access the Data Lake folder. Which three actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Register an Azure Active Directory app that uses the Web app/API application type.
- B. Configure the application to use the application ID, authentication key, and tenant ID.
- C. Assign the Azure Active Directory app permission to the Data Lake Store folder.
- D. Configure the application to use the OAuth 2.0 token endpoint.
- E. Register an Azure Active Directory app that uses the Native application type.
- F. Configure the application to use the application ID and redirect URI.

**Answer:** ABC

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-service-to-service-authenticate-using-active-directory>

**NEW QUESTION 58**

You need to use the Cognition.Vision.FaceDetector() function in U-SQL to analyze images. Which attribute can you detect by using the function?

- A. gender
- B. race
- C. weight
- D. hair color

Answer: A

**NEW QUESTION 60**

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

Answer: BC

**NEW QUESTION 62**

DRAG DROP

You are troubleshooting job performance and failure issues for Microsoft Azure Data Lake Analytics jobs.

You need to perform the following tasks:

Which tool should you use for each task? To answer, drag the appropriate tools to the correct tasks. Each tool 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.

Tools	Answer Area
Diagnostic logs	View the start time and the end time of queries: <input type="text" value="Tool"/>
Job Browser	Identify the job steps that have the highest number of write operations: <input type="text" value="Tool"/>
Vertex Execution View	

Answer:

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/data-lake-analytics/data-lake-analytics-data-lake-tools-view-jobs>

**NEW QUESTION 65**

HOTSPOT

You have a Microsoft Azure Data Lake Analytics service.

You have a file named Employee.tsv that contains data on employees. Employee.tsv contains seven columns named EmplId, Start, FirstName, LastName, Age, Department, and Title.

You need to create a Data Lake Analytics jobs to transform Employee.tsv, define a schema for the data, and output the data to a CSV file. The outputted data must contain only employees who are in the sales department. The Age column must allow NULL.

How should you complete the U-SQL code segment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

```

@Employee =
    EXTRACT EmpID      int,
            Start      DateTime,
            FirstName   string,
            LastName    string,
            Department  string,
            Title       string,
            Age         
                    
                    int
                    int?
                    string

    FROM "/Input/Data/Employee.tsv"
    USING 
        Extractors.Csv();
        Extractors.Tsv();
        Outputters.Csv();
        Outputters.Tsv();

@Rowset =
    SELECT EmpId, Start, FirstName, LastName, Age, Department, Title
    FROM @Employee
    WHERE Department == "Sales";

OUTPUT @Rowset
    TO "/Output/Employee.csv"
    USING 
        Extractors.Csv();
        Extractors.Tsv();
        Outputters.Csv();
        Outputters.Tsv();
    
```

**Answer:**

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/data-lake-analytics/data-lake-analytics-u-sql-get-started>

**NEW QUESTION 70**

You have a Microsoft Azure Data Lake Analytics service.

You need to provide a user with the ability to monitor Data Lake Analytics jobs. The solution must minimize the number of permissions assigned to the user. Which role should you assign to the user?

- A. Reader
- B. Owner
- C. Contributor
- D. Data Lake Analytics Developer

**Answer:** A

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/data-lake-analytics/data-lake-analytics-manage-use-portal>

**NEW QUESTION 74**

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.

**Answer:** B

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-statistics>

**NEW QUESTION 77**

You have a Microsoft Azure SQL data warehouse. The following statements are used to define file formats in the data warehouse.

```
CREATE EXTERNAL FILE FORMAT FileFormat_ORC
WITH (
    FORMAT_TYPE = ORC
, DATA_COMPRESSION = 'org.apache.hadoop.io.compress.SnappyCodec'
);
```

```
CREATE EXTERNAL FILE FORMAT FileFormat_PARQUET
WITH (
    FORMAT_TYPE = PARQUET
, DATA_COMPRESSION = 'org.apache.hadoop.io.compress.SnappyCodec'
);
```

You have an external PolyBase table named file\_factPowerMeasurement that uses the FileFormat\_ORC file format.

You need to change file\_factPowerMeasurement to use the FileFormat\_PARQUET file format. Which two statements should you execute? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. CREATE EXTERNAL TABLE
- B. ALTER TABLE
- C. CREATE EXTERNAL TABLE AS SELECT
- D. ALTER EXTERNAL DATA SOURCE
- E. DROP EXTERNAL TABLE

**Answer:** AE

**NEW QUESTION 81**

You have a Microsoft Azure SQL data warehouse that has 10 compute nodes.

You need to export 10 TB of data from a data warehouse table to several new flat files in Azure Blob storage. The solution must maximize the use of the available compute nodes.

What should you do?

- A. Use the bcp utility.
- B. Execute the CREATE EXTERNAL TABLE AS SELECT statement.
- C. Create a Microsoft SQL Server Integration Services (SSIS) package that has a data flow task.
- D. Create a Microsoft SQL Server Integration Services (SSIS) package that has an SSIS Azure Blob Storage task.

**Answer:** D

**NEW QUESTION 86**

You have a fact table named PowerUsage that has 10 billion rows. PowerUsage contains data about customer power usage during the last 12 months. The usage data is collected every minute. PowerUsage contains the columns configured as shown in the following table.

Column name	Data type	Nullable
MeasurementId	bigint	No
CustomerId	int	No
LocationNumber	int	No
MinuteOfMonth	int	No
MonthKey	int	No
Usage	int	Yes

LocationNumber has a default value of 1. The MinuteOfMonth column contains the relative minute within each month. The value resets at the beginning of each month.

A sample of the fact table data is shown in the following table.

Measurement Id	CustomerId	Location Number	MinuteOf Month	MonthKey	Usage
1	1	1	1	1	100
2	1	1	2	1	66
3	2	2	1	1	88
4	1	1	1	2	93
5	1	1	2	2	0
6	2	2	1	2	47
7	1	1	1	2	52
8	1	1	2	2	22

There is a related table named Customer that joins to the PowerUsage table on the CustomerId column. Sixty percent of the rows in PowerUsage are associated to less than 10 percent of the rows in Customer. Most queries do not require the use of the Customer table. Many queries select on a specific month.

You need to minimize how long it takes to find the records for a specific month. What should you do?

- A. Implement partitioning by using the MonthKey column
- B. Implement hash distribution by using the CustomerId column.
- C. Implement partitioning by using the CustomerId column
- D. Implement hash distribution by using the MonthKey column.
- E. Implement partitioning by using the MonthKey column
- F. Implement hash distribution by using the MeasurementId column.

- G. Implement partitioning by using the MinuteOfMonth column
- H. Implement hash distribution by using the MeasurementId column.

**Answer:** C

**NEW QUESTION 87**

DRAG DROP

You have an on-premises Microsoft SQL Server instance named Instance1 that contains a database named DB1.

You have a Data Management Gateway named Gateway1.

You plan to create a linked service in Azure Data Factory for DB1.

You need to connect to DB1 by using standard SQL Server Authentication. You must use a username of User1 and a password of P@\$w0rd89.

How should you complete the JSON code? TO answer, drag the appropriate values to the correct targets. Each value 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.

Values	Answer Area
<input style="border: 1px solid black;" type="text" value='"external": false'/>	<pre> {   "name": "DataSource1",   "properties":   {     "type": "OnPremisesSqlServer",     "typeProperties":     {       "connectionString":       "Data Source=Instance1";       Initial Catalog=Db1;       <input style="border: 1px solid black;" type="text" value="Value"/> ;       User ID=User1;       Password=P@\$rd89";       <input style="border: 1px solid black;" type="text" value="Value"/>     }   } } </pre>
<input style="border: 1px solid black;" type="text" value='"external": true'/>	
<input style="border: 1px solid black;" type="text" value='"gatewayName": "Gateway1"'/>	
<input style="border: 1px solid black;" type="text" value="Integrated Security= False"/>	
<input style="border: 1px solid black;" type="text" value="Integrated Security= True"/>	

**Answer:**

**Explanation:**

References:

<https://github.com/uglide/azure-content/blob/master/articles/data-factory/data-factory-move-data-between-onprem-and-cloud.md>

**NEW QUESTION 89**

You have an on-premises Microsoft SQL Server instance.

You plan to copy a table from the instance to a Microsoft Azure Storage account. You need to ensure that you can copy the table by using Azure Data Factory.

Which service should you deploy?

- A. an on-premises data gateway
- B. Azure Application Gateway
- C. Data Management Gateway
- D. a virtual network gateway

**Answer:** C

**NEW QUESTION 90**

You need to define an input dataset for a Microsoft Azure Data Factory pipeline.

Which properties should you include when you define the dataset?

- A. name, type, typeProperties, and availability
- B. name, typeProperties, structure, and availability
- C. name, policy, structure, and external
- D. name, type, policy, and structure

**Answer:** A

**Explanation:**

References:

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

**NEW QUESTION 94**

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.

Cmdlets	Answer Area
Set-AzureRmDataFactoryGateway	
New-AzureRmDataFactoryLinkedService	
New-AzureRmDataFactoryDataset	
New-AzureRmDataFactoryGateway	
Set-AzureRmDataFactory	
New-AzureRmDataFactoryHub	
New-AzureRmDataFactoryPipeline	

>
<
^
v

**Answer:**

**Explanation:**

References:

<https://docs.microsoft.com/en-us/powershell/module/azurermdatadf/new-azurermdatadfactorypipeline?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>

**NEW QUESTION 96**

HOTSPOT

You are creating a series of activities for a Microsoft Azure Data Factory. The first activity will copy an input dataset named Dataset1 to an output dataset named Dataset2. The second activity will copy a dataset named Dataset3 to an output dataset named Dataset4.

Dataset1 is located in Azure Table Storage. Dataset2 is located in Azure Blob storage. Dataset3 is located in an Azure Data Lake store. Dataset4 is located in an Azure SQL data warehouse.

You need to configure the inputs for the second activity. The solution must ensure that Dataset3 is copied after Dataset2 is created.

How should you complete the JSON code for the second activity? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

```

...
"type": "Copy",
"typeProperties": {
  "source": {
    "type": 
  },
  "sink": {
    "type": 
  },
  "writeBatchSize": 0,
  "writeBatchTimeout": "00:00:00"
}
},
"inputs": [
  { "name":  },
  { "name":  },
  { "name":  },
  { "name":  }
]

```

**Answer:**

**Explanation:**

References:

<https://github.com/aelij/azure-content/blob/master/articles/data-factory/data-factory-create-pipelines.md>

**NEW QUESTION 98**

DRAG DROP

You need to create a dataset in Microsoft Azure Data Factory that meets the following requirements: How should you complete the JSON code? To answer, drag

the appropriate values to the correct targets. Each value 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.

Values	Answer Area
<input type="text" value="availability"/>	<pre> {   "name": "blob1",   "properties": {     "type": "AzureBlob",     "linkedServiceName": "LinkedService1",     "typeProperties": {       "folderPath": "Container1/myfolder/{Year}/{Month}"       <input type="text" value="Value"/>       [         {           "name": "Year",           "value": {             "type": "DateTime",             "date": "SliceStart",             "format": "yyyy"           }         },         {           "name": "Month",           "value": {             "type": "DateTime",             "date": "SliceStart",             "format": "MM"           }         }       ]     }   },   <input type="text" value="Value"/>   {     "frequency": "Month",     "interval": 1   } } </pre>
<input type="text" value="partitionedBy"/>	
<input type="text" value="policy"/>	
<input type="text" value="scheduler"/>	

**Answer:**

**Explanation:**

References:  
<https://github.com/aelij/azure-content/blob/master/articles/data-factory/data-factory-create-pipelines.md>

**NEW QUESTION 101**

DRAG DROP

You plan to create for an alert for a Microsoft Azure Data Factory pipeline. You need to configure the alert to trigger when the total number of failed runs exceeds five within a three-hour period. How should you configure the window size and the threshold in the JSON file? To answer, drag the appropriate values to the correct targets. Each value 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.

Values	Answer Area
<input type="text" value="3.0"/>	<p>threshold: <input type="text" value="Value"/></p> <p>windowSize: <input type="text" value="Value"/></p>
<input type="text" value="5.0"/>	
<input type="text" value="PT3H"/>	
<input type="text" value="PT5H"/>	

**Answer:**

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/data-factory/v1/data-factory-monitor-manage-pipelines?view=powerbiapi-1.1.10>

**NEW QUESTION 103**

You have a Microsoft Azure SQL data warehouse named DW1 that is used only from Monday to Friday. You need to minimize Data Warehouse Unit (DWU) usage during the weekend. What should you do?

- A. From the Azure CLI, run the account set command.
- B. Run the ALTER DATABASE statement.
- C. Call the Create or Update Database REST API.
- D. Run the Suspend-AzureRmSqlDatabase Azure PowerShell cmdlet.

**Answer:** D

**NEW QUESTION 108**

You have a Microsoft Azure Data Lake Store and an Azure Active Directory tenant.

You are developing an application that will access the Data Lake Store by using end-user credentials. You need to ensure that the application uses end-user authentication to access the Data Lake Store. What should you create?

- A. a Native Active Directory app registration
- B. a policy assignment that uses the Allowed resource types policy definition
- C. a Web app/API Active Directory app registration
- D. a policy assignment that uses the Allowed locations policy definition

**Answer:** A

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-end-user-authenticate-using-active-directory>

**NEW QUESTION 113**

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.

**Answer Area**

<p><b>Job1:</b></p> <div style="border: 1px solid black; padding: 5px; display: inline-block;"> <input style="width: 100%; height: 20px;" type="text"/> ▼         </div> <div style="border: 1px solid black; padding: 5px; margin-top: 5px;"> <p>1</p> <p>3</p> <p>6</p> <p>12</p> <p>18</p> <p>24</p> <p>30</p> <p>32</p> <p>42</p> <p>48</p> </div>	<p><b>Job2:</b></p> <div style="border: 1px solid black; padding: 5px; display: inline-block;"> <input style="width: 100%; height: 20px;" type="text"/> ▼         </div> <div style="border: 1px solid black; padding: 5px; margin-top: 5px;"> <p>1</p> <p>3</p> <p>6</p> <p>12</p> <p>18</p> <p>24</p> <p>30</p> <p>32</p> <p>42</p> <p>48</p> </div>
--	--

**Answer:**

**Explanation:**

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>

**NEW QUESTION 117**

DRAG DROP

You have a Microsoft Azure SQL data warehouse.

Users discover that reports running in the data warehouse take longer than expected to complete. You need to review the duration of the queries and which users are running the queries currently. Which dynamic management view should you review for each requirement? To answer, drag the appropriate dynamic management views to the correct requirements. Each dynamic management view 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.

Dynamic Management Views	Answer Area
<input type="text" value="Sys.dm_pdw_exec_requests"/>	Duration of the queries: <input type="text" value="Dynamic Management Views"/>
<input type="text" value="Sys.dm_pdw_exec_sessions"/>	Which users are running queries currently: <input type="text" value="Dynamic Management Views"/>
<input type="text" value="Sys.dm_pdw_os_threads"/>	
<input type="text" value="Sys.dm_pdw_request_steps"/>	

**Answer:**

**Explanation:**

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-exec-requests-transact-sql>

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-exec-sessions-transact-sql>

**NEW QUESTION 122**

DRAG DROP

You have a Microsoft Azure SQL data warehouse named DW1. Data is loaded to DW1 once daily at 01:00.

A user accidentally deletes data from a fact table in DW1 at 09:00.

You need to recover the lost data. The solution must prevent the need to change any connection strings and must minimize downtime.

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

Actions	Answer Area
<input type="text" value="Resume DW1."/>	
<input type="text" value="Restore the database to a database named DW2."/>	
<input type="text" value="Pause DW2."/>	
<input type="text" value="Delete DW1."/>	
<input type="text" value="Rename DW2."/>	
<input type="text" value="Restore the database to a database named DW1."/>	
<input type="text" value="Pause DW1."/>	

**Answer:**

**Explanation:**

Actions	Answer Area
<input type="text" value="Resume DW1."/>	<input type="text" value="Restore the database to a database named DW2."/>
<input type="text" value="Restore the database to a database named DW2."/>	<input type="text" value="Delete DW1."/>
<input type="text" value="Pause DW2."/>	<input type="text" value="Rename DW2."/>
<input type="text" value="Delete DW1."/>	
<input type="text" value="Rename DW2."/>	
<input type="text" value="Restore the database to a database named DW1."/>	
<input type="text" value="Pause DW1."/>	

**NEW QUESTION 123**

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

**Answer:** E

**NEW QUESTION 128**

You manage an on-premises data warehouse that uses Microsoft SQL Server. The data warehouse contains 100 TB of data. The data is partitioned by month. One TB of data is added to the data warehouse each month.

You create a Microsoft Azure SQL data warehouse and copy the on-premises data to the data warehouse.

You need to implement a process to replicate the on-premises data warehouse to the Azure SQL data warehouse. The solution must support daily incremental updates and must provide error handling. What should you use?

- A. the Azure Import/Export service
- B. SQL Server log shipping
- C. Azure Data Factory
- D. the AzCopy utility

**Answer: C**

**NEW QUESTION 131**

DRAG DROP

You need to load data from Microsoft Azure Data Lake Store to Azure SQL Data Warehouse by using Transact-SQL.

In which sequence 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.

Actions	Answer Area
Use the CREATE TABLE AS SELECT feature.	
Create a credential.	
Create a data format.	
Create the external data source.	
Create external tables.	

**Answer:**

**Explanation: Actions**

Answer Area
1 Create a credential.
2 Create the external data source.
3 Create a data format.
4 Create external tables.
5 Use the CREATE TABLE AS SELECT feature.

**NEW QUESTION 136**

You are designing a solution that will use Microsoft Azure Data Lake Store.

You need to recommend a solution to ensure that the storage service is available if a regional outage occurs. The solution must minimize costs. What should you recommend?

- A. Create two Data Lake Store accounts and copy the data by using Azure Data Factory.
- B. Create one Data Lake Store account that uses a monthly commitment package.
- C. Create one read-access geo-redundant storage (RA-GRS) account and configure a Recovery Services vault.
- D. Create one Data Lake Store account and create an Azure Resource Manager template that redeploys the services to a different region.

**Answer: D**

**NEW QUESTION 139**

HOTSPOT

You have a Microsoft Azure Data Factory version 2 (V2) service.

You need to deploy a linked service to an existing database in Azure SQL Data Warehouse.

How should you complete the JSON snippet? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```

{
  "name": "LinkedService1",
  "properties": {
    "type": [
      "AzureDataLakeAnalytics",
      "AzureSqlDatabase",
      "AzureSqlDW",
      "AzureStorage"
    ],
    "typeProperties": {
      "connection": {
        "connectionString": {
          "name": {
            "properties": {
              "value": {
                "type": "SecureString",
                "value": "Integrated Security=False;
                Encrypt=True;
                Connection Timeout=30;
                Data Source=Adatum-02.database.windows.net;
                Initial Catalog=AdatumCatalog01;
                User ID=AdatumUser01;
                Password=PA$$w0rd001!"
              }
            }
          }
        }
      }
    }
  }
}

```

Answer:

Explanation: Answer Area

```

"properties": {
  "type": "AzureSqlDW",
  "typeProperties": {
    "connection": {
      "type": "SecureString",
      "value": "Integrated Security=False;
      Encrypt=True;
      Connection Timeout=30;
    }
  }
}

```

**NEW QUESTION 142**

DRAG DROP

You need to create a linked service in Microsoft Azure Data Factory. The linked service must use an Azure Database for MySQL table named Customers. How should you complete the JSON snippet? To answer, drag the appropriate values to the correct targets. Each value 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.

**Values**

- 
- 
- 
- 
- 
-

**Answer Area**

```

{
  "name": "AzureMySQLDataset",
  "properties": {
    "type": [ Value ],
    "linkedServiceName": {
      "referenceName": "mysql.database.windows.net",
      "type": [ Value ]
    }
  }
}

```

Answer:

Explanation: Values

- 
- 
- 
- 
- 
- 

Answer Area

```

{
  "name": "AzureMySQLDataset",
  "properties": {
    "type": "AzureMySQLTable",
    "linkedServiceName": {
      "referenceName": "mysql.database.windows.net",
      "type": "MySQL"
    }
  }
}

```

**NEW QUESTION 143**

You are implementing a solution by using Microsoft Azure Data Lake Analytics. You have a dataset that contains data-related to website visits. You need to combine overlapping visits into a single entry based on the timestamp of the visits. Which type of U-SQL interface should you use?

- A. IExtractor
- B. IReducer
- C. Aggregate
- D. IProcessor

**Answer: C**

**NEW QUESTION 145**

You have a Microsoft Azure Data Lake Store that contains a folder named /Users/User1 and an Azure Active Directory account named User1.

You need to provide access to the Data Lake Store to meet the following requirements:

- Grant User1 read and list access to /Users/User1.
- Prevent User1 from viewing the contents in /Users.
- Minimize the number of permissions granted to User1. What should you do?

- A. Grant User1 Execute permissions to /Users and /Users/User1.
- B. Grant User1 Read permissions to /Users folder and /Users/User1.
- C. Grant User1 Read permissions to Users/User1.
- D. Grant User1 Execute permissions to /User
- E. Grant User1 Read & Execute permissions to /Users/User1.

**Answer: D**

**NEW QUESTION 146**

You plan to use Microsoft Azure Event Hubs in Azure Stream Analytics to consume time-series aggregations from several published data sources, such as IoT data, reference data, and social media. You expect several TB of data to be consumed daily. All the consumed data will be retained for one week.

You need to recommend a storage solution for the data. The solution must minimize costs. What should you recommend?

- A. Azure DocumentDB
- B. Azure Data Lake
- C. Azure Table Storage
- D. Azure Blob storage

**Answer: B**

**NEW QUESTION 151**

You have the following process:

- A CSV file is ingested by Microsoft Azure Stream Analytics.
- Scoring is performed by Azure Machine Learning.
- Stream Analytics returns sentiment scoring through a web service endpoint.
- Stream Analytics creates an output blob.

You need to view the output of the scoring operation and to evaluate the throughput to the Machine Learning models.

Which monitoring data should you evaluate from the Azure portal?

- A. the request count of Stream Analytics
- B. the request count of Machine Learning
- C. the event count of Stream Analytics
- D. the event count of Machine Learning

**Answer: C**

**NEW QUESTION 155**

You have a Microsoft Azure Stream Analytics job.

You are debugging event information manually.

You need to view the event data that is being collected.

Which monitoring data should you view for the Stream Analytics job?

- A. query
- B. outputs
- C. scale
- D. inputs

**Answer: D**

**NEW QUESTION 158**

DRAG DROP

You are building a data pipeline that uses Microsoft Azure Stream Analytics.

Alerts are generated when the aggregate of data streaming in from devices during a minute-long window matches the values in a rule.

You need to retrieve the following information:

\*The event ID



**Answer Area**

```
CREATE DATABASE dw1  
(  
    DISTRIBUTED_SIZE = 10485760,  
    Edition = 'Datawarehouse',  
    SERVICE_OBJECTIVE=DW1000  
)
```

**NEW QUESTION 169**

You have a Microsoft Azure SQL data warehouse that has 10 compute nodes.

You need to export 10 TB of data from a data warehouse table to several new flat files in Azure Blob storage. The solution must maximize the use of the available compute nodes.

What should you do?

- A. Use a transform activity in an Azure Data Factory pipeline.
- B. Execute the create table as select statement.
- C. Use the bcp utility.
- D. Execute the create external table as select statement.

**Answer: C**

**NEW QUESTION 171**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your 70-776 Exam with Our Prep Materials Via below:**

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