

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 nonclustered 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 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

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 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

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 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 7

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 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.

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 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. 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 10

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

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 Area

Answer:

Explanation:

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

NEW QUESTION 12

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 13

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 14

You have sensor devices that report data to Microsoft Azure Stream Analytics. Each sensor reports data several times per second.

You need to create a live dashboard in Microsoft Power BI that shows the performance of the sensor devices. The solution must minimize lag when visualizing the data.

Which function should you use for the time-series data element?

- A. LAG
- B. SlidingWindow
- C. System.TimeStamp
- D. TumblingWindow

Answer: D

NEW QUESTION 17

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 using a Microsoft Azure Stream Analytics query language. You are outputting data from an input click stream. You need to ensure that when you consecutively receive two rows from the same IP address within one minute, only the first row is outputted. Which functions should you use in the WHERE statement?

- A. Last and HoppingWindow
- B. Last and SlidingWindow
- C. LAG and HoppingWindow
- D. LAG and Duration

Answer: B

NEW QUESTION 23

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 28

You plan to use Microsoft Azure Event Hubs to ingest sensor data. You plan to use Azure Stream Analytics to analyze the data in real time and to send the output directly to Azure Data Lake Store. You need to write events to the Data Lake Store in batches. What should you use?

- A. Apache Storm in Azure HDInsight
- B. Stream Analytics
- C. Microsoft SQL Server Integration Services (SSIS)
- D. the Azure CLI

Answer: B

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-data-scenarios>

NEW QUESTION 32

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

Answer:

Explanation:

References:

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

NEW QUESTION 36

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 40

You plan to use Microsoft Azure Event Hubs to ingest data. You plan to use Azure Stream Analytics to analyze the data in real time and to send the output directly to Azure Data Lake Store.

You discover duplicate records in the output data. What is a possible cause of the duplicate records?

- A. There are connectivity issues with the output adapter.
- B. There is a connectivity issue between the data source and the event hub.
- C. There are multiple deliveries to the output adapter that writes the output events.
- D. The Stream Analytics output adapter writes the output events transactionally.

Answer: A

Explanation:

References:

<https://msdn.microsoft.com/en-us/library/azure/mt721300.aspx>

NEW QUESTION 45

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 47

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 52

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 56

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 57

You are developing an application by using the Microsoft .NET SDK. The application will access data from a Microsoft Azure Data Lake folder.

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

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

Answer Area

```
CREATE TABLE dbo.UserActivity
(
    INDEX s_idx CLUSTERED (SessionId ASC)
    DISTRIBUTED BY  (SessionId) INTO 2
    HASH
    RANGE
    ROUND ROBIN
)
EXTRACT SessionId    Guid
      , Timestamp    DateTime
      , Geocode       string
      , ResourcePath  string
      , ExecutionTime decimal
      , HeaderLog     string
FROM "/data/UserActivity.tsv"
USING 
;
```

<https://msdn.microsoft.com/en-us/library/mt706197.aspx>

NOTE: Each correct selection is worth one point.

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

Answer:

Explanation:

References:

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

NEW QUESTION 72

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 EmpId, 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 74

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

Answer: B

Explanation:

References:

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

NEW QUESTION 75

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.

Answer Area

Create a U-SQL project:	<div style="border: 1px solid black; padding: 2px;"> <div style="border-bottom: 1px solid black; height: 1.2em; margin-bottom: 2px;"></div> <div style="display: flex; justify-content: flex-end; align-items: center;">▼</div> <div style="border: 1px solid black; padding: 2px; margin-top: 2px;"> Azure portal only Visual Studio only Both Visual Studio and Azure portal </div> </div>
Create an Azure Data Lake Analytics account:	<div style="border: 1px solid black; padding: 2px;"> <div style="border-bottom: 1px solid black; height: 1.2em; margin-bottom: 2px;"></div> <div style="display: flex; justify-content: flex-end; align-items: center;">▼</div> <div style="border: 1px solid black; padding: 2px; margin-top: 2px;"> Azure portal only Visual Studio only Both Visual Studio and Azure portal </div> </div>
Upload data files to the default Azure Data Lake Store:	<div style="border: 1px solid black; padding: 2px;"> <div style="border-bottom: 1px solid black; height: 1.2em; margin-bottom: 2px;"></div> <div style="display: flex; justify-content: flex-end; align-items: center;">▼</div> <div style="border: 1px solid black; padding: 2px; margin-top: 2px;"> Azure portal only Visual Studio only Both Visual Studio and Azure portal </div> </div>

Answer:

Explanation: **Answer Area**

Create a U-SQL project:	<div style="border: 1px solid black; padding: 2px;"> <div style="border-bottom: 1px solid black; height: 1.2em; margin-bottom: 2px;"></div> <div style="display: flex; justify-content: flex-end; align-items: center;">▼</div> <div style="border: 1px solid black; padding: 2px; margin-top: 2px;"> Azure portal only Visual Studio only Both Visual Studio and Azure portal </div> </div>
Create an Azure Data Lake Analytics account:	<div style="border: 1px solid black; padding: 2px;"> <div style="border-bottom: 1px solid black; height: 1.2em; margin-bottom: 2px;"></div> <div style="display: flex; justify-content: flex-end; align-items: center;">▼</div> <div style="border: 1px solid black; padding: 2px; margin-top: 2px;"> Azure portal only Visual Studio only Both Visual Studio and Azure portal </div> </div>
Upload data files to the default Azure Data Lake Store:	<div style="border: 1px solid black; padding: 2px;"> <div style="border-bottom: 1px solid black; height: 1.2em; margin-bottom: 2px;"></div> <div style="display: flex; justify-content: flex-end; align-items: center;">▼</div> <div style="border: 1px solid black; padding: 2px; margin-top: 2px;"> Azure portal only Visual Studio only Both Visual Studio and Azure portal </div> </div>

NEW QUESTION 76

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 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 83

HOTSPOT

You are designing a fact table that has 100 million rows and 1,800 partitions. The partitions are defined based on a column named OrderDayKey. The fact table will contain:

Data from the last five years

A clustered columnstore index

A column named YearMonthKey that stores the year and the month

Multiple transformations will be performed on the fact table during the loading process. The fact table will be hash distributed on a column named OrderId.

You plan to load the data to a staging table and to perform transformations on the staging table. You will then load the data from the staging table to the final fact table.

You need to design a solution to load the data to the fact table. The solution must minimize how long it takes to perform the following tasks:

Load the staging table.

Transfer the data from the staging table to the fact table. Remove data that is older than five years.

Query the data in the fact table

How should you configure the tables? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

Configure the table as a heap:

▼

☐ Fact table only
☐ Staging table only
☐ Staging and fact tables

Configure the table to use a clustered columnstore index:

▼

☐ Fact table only
☐ Staging table only
☐ Staging and fact tables

Partition the table by YearMonthKey:

▼

☐ Fact table only
☐ Staging table only
☐ Staging and fact tables

Answer:

Explanation:

Answer Area

Configure the table as a heap:

Fact table only

Staging table only

Staging and fact tables

Configure the table to use a clustered columnstore index:

Fact table only

Staging table only

Staging and fact tables

Partition the table by YearMonthKey:

Fact table only

Staging table only

Staging and fact tables

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

"external": false

"external": true

"gatewayName": "Gateway1"

Integrated Security= False

Integrated Security= True

Answer Area

```
{
  "name": "DataSource1",
  "properties": {
    "type": "OnPremisesSqlServer",
    "typeProperties": {
      "connectionString":
        "Data Source=Instance1;
        Initial Catalog=Db1;
        Value
        ;
        User ID=User1;
        Password=P@$rd89;",
        Value
      }
    }
  }
}
```

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 92

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 97

You have a Microsoft Azure Data Factory that recently ran several activities in parallel. You receive alerts indicating that there are insufficient resources.

From the Activity Windows list in the Monitoring and Management app, you discover the statuses described in the following table.

Activity name	Status	Substatus
Activity1	Failed	Canceled
Activity2	Waiting	DatasetDependencies
Activity3	Waiting	ComputeResources
Activity4	Waiting	ActivityResume
Activity5	Waiting	ConcurrencyLimit
Activity6	Skipped	Not applicable
Activity7	In progress	Validating
Activity8	Waiting	ValidationRetry

Which activity cannot complete because of insufficient resources?

- A. Activity2
- B. Activity4
- C. Activity5
- D. Activity7

Answer: C

NEW QUESTION 100

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":  }
  ]
}

```

Available options for source and sink type dropdowns:

- AzureDataLakeStore
- AzureSqlDatabase
- AzureSqlDW
- AzureStorage

Available options for input dataset name dropdowns:

- Dataset1
- Dataset2
- Dataset3
- Dataset4

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 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

"availability":

"partitionedBy":

"policy":

"scheduler":

Answer Area

```
{
  "name": "blob1",
  "properties": {
    "type": "AzureBlob",
    "linkedServiceName": "LinkedService1",
    "typeProperties": {
      "folderPath": "Container1/myfolder/{Year}/{Month}"
    }
  }
}
```

Value

```
[
  {
    "name": "Year",
    "value": {
      "type": "DateTime",
      "date": "SliceStart",
      "format": "yyyy"
    }
  },
  {
    "name": "Month",
    "value": {
      "type": "DateTime",
      "date": "SliceStart",
      "format": "MM"
    }
  }
]
},
{
  "frequency": "Month",
  "interval": 1
}
}
```

Value

Answer:

Explanation:

References:

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

NEW QUESTION 102

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

3.0

5.0

"PT3H"

"PT5H"

Answer Area

threshold:

Value

windowSize:

Value

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 104

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

Job1:	<div><div></div><div>▼</div></div>	Job2:	<div><div></div><div>▼</div></div>
1		1	
3		3	
6		6	
12		12	
18		18	
24		24	
30		30	
32		32	
42		42	
48		48	

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 109

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
<div>Sys.dm_pdw_exec_requests</div>	Duration of the queries: <div>Dynamic Management Views</div>
<div>Sys.dm_pdw_exec_sessions</div>	Which users are running queries currently: <div>Dynamic Management Views</div>
<div>Sys.dm_pdw_os_threads</div>	
<div>Sys.dm_pdw_request_steps</div>	

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 113

DRAG DROP

You are monitoring a Microsoft Azure SQL data warehouse. You need to get the following information:

- The top 10 longest running queries
- The distributed query plan for a specific query

Which dynamic management view should you use for each piece of information? To answer, drag the appropriate dynamic management views to the correct pieces of information. 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
sys.dm_pdw_dms_workers	
sys.dm_pdw_exec_requests	The top 10 longest running queries: Dynamic Management Views
sys.dm_pdw_exec_sessions	The distributed query plan for a specific query: Dynamic Management Views
sys.dm_pdw_request_steps	
sys.dm_pdw_waits	

Answer:

Explanation:

References:

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

NEW QUESTION 115

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
Resume DW1.	
Restore the database to a database named DW2.	
Pause DW2.	
Delete DW1.	
Rename DW2.	
Restore the database to a database named DW1.	
Pause DW1.	

Answer:

Explanation:

Actions	Answer Area
Resume DW1.	Restore the database to a database named DW2.
Restore the database to a database named DW2.	Delete DW1.
Pause DW2.	Rename DW2.
Delete DW1.	
Rename DW2.	
Restore the database to a database named DW1.	
Pause DW1.	

NEW QUESTION 118

You have an on-premises data warehouse that uses Microsoft SQL Server 2016. All the data in the data warehouse comes from text files stored in Azure Blob storage. The text files are imported into the data warehouse by using SQL Server Integration Services (SSIS). The text files are not transformed.

You need to migrate the data to an Azure SQL data warehouse in the least amount of time possible. Which two actions should you perform? Each correct answer

presents part of the solution.
NOTE: Each correct selection is worth one point.

- A. Use SSIS to upload the files in Azure Blob storage to tables in the Azure SQL data warehouse.
- B. Execute the CREATE EXTERNAL TABLE AS SELECT statement to export the data.
- C. Use AzCopy to transfer the data from the on-premises data warehouse to Azure SQL data warehouse.
- D. Execute the CREATE TABLE AS SELECT statement to load the data.
- E. Define external tables in the Azure SQL data warehouse that map to the existing files in Azure Blob storage.

Answer: DE

Explanation:

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

NEW QUESTION 123

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

Use the CREATE TABLE AS SELECT feature.

Create a credential.

Create a data format.

Create the external data source.

Create external tables.

Answer Area

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 128

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 133

DRAG DROP

You are designing a Microsoft Azure analytics solution. The solution requires that data be copied from Azure Blob storage to Azure Data Lake Store.
The data will be copied on a recurring schedule. Occasionally, the data will be copied manually. You need to recommend a solution to copy the data.
Which tools should you include in the recommendation? To answer, drag the appropriate tools to the correct requirements. 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

AdlCopy

AzCopy

Azure Data Factory

distcp

sqoop

Answer Area

To manually copy the data:

Tool

To schedule the copying of the data:

Tool

Answer:

Explanation: Tools

AdlCopy

AzCopy

Azure Data Factory

distcp

sqoop

Answer Area

To manually copy the data: AdlCopy

To schedule the copying of the data: Azure Data Factory

NEW QUESTION 136

DRAG DROP

You plan to use U-SQL to run federated queries to join data from a Microsoft Azure SQL data warehouse, an Azure SQL database, and a Microsoft SQL Server database on an Azure virtual machine.

You need to ensure that you can use a U-SQL query that joins all three data sources. The solution must ensure that the three data sources appear as tables in U-SQL without having to move the data from external sources.

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

Statements

CREATE TABLE

CREATE DATA SOURCE

CREATE CREDENTIAL

CREATE EXTERNAL TABLE

CREATE PACKAGE

Answer Area

>

<

⬆

⬆

Answer:

Explanation: Statements

CREATE TABLE

CREATE PACKAGE

Answer Area

1 CREATE CREDENTIAL

2 CREATE DATA SOURCE

3 CREATE EXTERNAL TABLE

⬆

⬆

NEW QUESTION 137

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 138

You have a Microsoft Azure Data Lake Analytics service.

You need to store a list of milltiple-character string values in a single column and to use a cross apply explode expression to output the values.

Which type of data should you use in a U-SQL query?

- A. SQL.MAP
- B. SQL.ARRAY
- C. string
- D. byte []

Answer: B

NEW QUESTION 142

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 147

You plan to capture the output from a group of 500 IoT devices that produce approximately 10 GB of data per hour by using Microsoft Azure Stream Analytics. The data will be retained for one year. Once the data is processed, it will be stored in Azure, and then analyzed by using an Azure HDInsight cluster. You need to select where to store the output data from Stream Analytics. The solution must minimize costs. What should you select?

- A. Azure Table Storage
- B. Azure SQL Database
- C. Azure Blob storage
- D. Azure SQL Data Warehouse

Answer: C

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 152

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 155

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
- *The device ID
- *The application ID that runs the service

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

For each event key, store the threshold values for the rule in the reference data.

Join the events to the reference data by using the event key.

Join the cache ID from the reference data to the event ID.

Use the threshold values as the event ID.

Use the lookup value from the reference data as the threshold value.

Answer Area

Answer:

Explanation:

Actions		Answer Area
Join the cache ID from the reference data to the event ID.	<div>></div> <div><</div>	1 For each event key, store the threshold values for the rule in the reference data.
Use the threshold values as the event ID.		2 Join the events to the reference data by using the event key.
		3 Use the lookup value from the reference data as the threshold value.

NEW QUESTION 156

You plan to deploy a Microsoft Azure Stream Analytics job to filter multiple input streams from IoT devices that have a total data flow of 30 MB/s. You need to calculate how many streaming units you require for the job. The solution must prevent lag. What is the minimum number of streaming units required?

- A. 3
- B. 10
- C. 30
- D. 300

Answer: C

NEW QUESTION 161

HOTSPOT

You need to create a Microsoft Azure SQL data warehouse named dw1 that supports up to 10 TB of data. How should you complete the statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
CREATE DATABASE dw1
(
  
```

DISTRIBUTED_SIZE = 10240,
 DISTRIBUTED_SIZE = 10485760,
 MAXSIZE = 10,
 MAXSIZE = 10240,
 SIZE = 10,
 SIZE = 10240,

Edition = 'Basic',
 Edition = 'Datawarehouse',
 Edition = 'Premium',
 Edition = 'PremiumRS',
 Edition = 'Standard',

```

  SERVICE_OBJECTIVE=DW1000
)

```

Answer:

Explanation: Answer Area

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

```

NEW QUESTION 165

Note: This question is part of a series of questions that use 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.

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

Answer: A

NEW QUESTION 170

DRAG DROP

You have a Microsoft Azure SQL data warehouse.

You plan to reference data from Azure Blob storage. The data is stored in the GZIP compressed format. The blob storage requires authentication.

You create a master key for the data warehouse and a database schema.

You need to reference the data without importing the data to the data warehouse.

Which four statements should you execute in sequence? To answer, move the appropriate statements from the list of statements to the answer area and arrange them in the correct order.

Statements

CREATE EXTERNAL TABLE

CREATE TABLE AS SELECT

CREATE EXTERNAL DATA SOURCE

CREATE EXTERNAL FILE FORMAT

CREATE DATABASE SCOPED CREDENTIAL

>

<

Answer Area

<

>

Answer:

Explanation: Statements

CREATE TABLE AS SELECT

>

<

Answer Area

1 CREATE DATABASE SCOPED CREDENTIAL

2 CREATE EXTERNAL DATA SOURCE

3 CREATE EXTERNAL FILE FORMAT

4 CREATE EXTERNAL TABLE

<

>

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>