

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 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 implement round robin for table distribution.

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 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 5

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 6

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 7

DRAG DROP

You have IoT devices that produce the following output.

```
sourcecode language='javascript' padlinenumbers='true' ]
[
  {
    "devID": "8656787",
    "timestamp": "2017-05-31T10:21:00",
    "readings": [
      {
        "type": "SensorA",
        "value": 18.965
      },
      {
        "type": "SensorB",
        "value": 72.9157
      },
      {
        "type": "SensorC",
        "value": 1524.672
      }
    ]
  }
]
[/sourcecode]
```

You need to use Microsoft Azure Stream Analytics to convert the output into the tabular format described in the following table.

Timestamp	DevId	SensorA	SensorB	SensorC
2017-05:31T10:21:00Z	8656787	18.965	72.9157	1524.672

How should you complete the Stream Analytics query? 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
<div style="border: 1px solid black; padding: 2px; margin-bottom: 2px;">Cross Apply</div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 2px;">GetArrayElement</div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 2px;">GetArrayElements</div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 2px;">Outer Apply</div>	<pre>[sourcecode language='sql'] SELECT i.arrayvalue.timestamp, i.arrayvalue.devId, udf.getValue('SensorA', i.arrayvalue.readings) as SensorA, udf.getValue('SensorB', i.arrayvalue.readings) as SensorB, udf.getValue('SensorC', i.arrayvalue.readings) as SensorC FROM input Value Value (input.devices) as i [/sourcecode]</pre>

Answer:

Explanation:

Values	Answer Area
<input type="text" value="Cross Apply"/>	<pre>[sourcecode language='sql'] SELECT i.arrayvalue.timestamp, i.arrayvalue.devid, udf.getValue('SensorA', i.arrayvalue.readings) as SensorA, udf.getValue('SensorB', i.arrayvalue.readings) as SensorB, udf.getValue('SensorC', i.arrayvalue.readings) as SensorC FROM input <input type="text" value="Cross Apply"/> <input type="text" value="GetArrayElements"/> (input.devices) as i [/sourcecode]</pre>
<input type="text" value="GetArrayElement"/>	
<input type="text" value="GetArrayElements"/>	
<input type="text" value="Outer Apply"/>	

NEW QUESTION 8

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 9

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 10

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 10

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 11

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 15

HOTSPOT

You have a Microsoft Azure Data Lake Analytics service.

You have a tab-delimited file named UserActivity.tsv that contains logs of user sessions. The file does not have a header row.

You need to create a table and to load the logs to the table. The solution must distribute the data by a column named SessionId.

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

NOTE: Each correct selection is worth one point.

Answer Area

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

Answer:

Explanation:

References:

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

NEW QUESTION 19

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 22

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 25

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:

Create an Azure Data Lake Analytics account:

Upload data files to the default Azure Data Lake Store:

Answer:

Explanation: Answer Area

Create a U-SQL project:

Create an Azure Data Lake Analytics account:

Upload data files to the default Azure Data Lake Store:

NEW QUESTION 28

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 29

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 33

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 37

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 41

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
"external": false	<pre> { "name": "DataSource1", "properties": { "type": "OnPremisesSqlServer", "typeProperties": { "connectionString": "Data Source=Instance1; Initial Catalog=Db1; Value ; User ID=User1; Password=P@\$rd89;", Value } } } </pre>
"external": true	
"gatewayName": "Gateway1"	
Integrated Security= False	
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 45

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 49

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 51

DRAG DROP

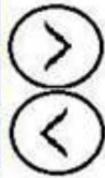
You have an Apache Hive database in a Microsoft Azure HDInsight cluster. You create an Azure Data Factory named DF1.

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



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 53

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 56

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}" } } } </pre>
<input type="text" value="partitionedBy"/>	
<input type="text" value="policy"/>	
<input type="text" value="scheduler"/>	
	<div style="border: 1px solid black; padding: 2px; width: fit-content; margin: 5px auto;">Value</div> <pre> [{ "name": "Year", "value": { "type": "DateTime", "date": "SliceStart", "format": "yyyy" } }, { "name": "Month", "value": { "type": "DateTime", "date": "SliceStart", "format": "MM" } }] </pre>
	<div style="border: 1px solid black; padding: 2px; width: fit-content; margin: 5px auto;">Value</div> <pre> { "frequency": "Month", "interval": 1 } </pre>

Answer:

Explanation:

References:

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

NEW QUESTION 61

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"/>	threshold: <input type="text" value="Value"/>
<input type="text" value="5.0"/>	windowSize: <input type="text" value="Value"/>
<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 62

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 65

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 67

You have an on-premises deployment of Active Directory named contoso.com. You plan to deploy a Microsoft Azure SQL data warehouse.

You need to ensure that the data warehouse can be accessed by contoso.com users.

Which two components should you deploy? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Azure AD Privileged Identity Management
- B. Azure Information Protection
- C. Azure Active Directory
- D. Azure AD Connect
- E. Cloud App Discovery
- F. Azure Active Directory B2C

Answer: CD

NEW QUESTION 72

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 74

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 77

You use Microsoft Azure Data Lake Store as the default storage for an Azure HDInsight cluster. You establish an SSH connection to the HDInsight cluster. You need to copy files from the HDInsight cluster to the Data LakeStore. Which command should you use?

- A. AzCopy
- B. hdfs dfs
- C. hadoop fs
- D. AdlCopy

Answer: D

NEW QUESTION 80

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 84

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
"AzureMySQLTable"	{
"AzureSQLDatabase"	"name": "AzureMySQLDataset",
"AzureStorage"	"properties": {
"DataSetReference"	"type": Value
"LinkedServiceReference"	"linkedServiceName": {
"MySQL"	"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 89

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

-
-
-
-
-

Answer Area

To manually copy the data:

To schedule the copying of the data:

Answer:

Explanation: Tools

-
-
-
-
-

Answer Area

To manually copy the data:

To schedule the copying of the data:

NEW QUESTION 90

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 94

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 95

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 99

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 104

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 106

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

- Join the cache ID from the reference data to the event ID.
- Use the threshold values as the event ID.

Answer Area

- 1 For each event key, store the threshold values for the rule in the reference data.
- 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 111

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		Answer Area
CREATE EXTERNAL TABLE		
CREATE TABLE AS SELECT	➤	⬆
CREATE EXTERNAL DATA SOURCE	⬅	⬇
CREATE EXTERNAL FILE FORMAT		
CREATE DATABASE SCOPED CREDENTIAL		

Answer:

Explanation:

Statements		Answer Area
CREATE TABLE AS SELECT	➤	1 CREATE DATABASE SCOPED CREDENTIAL
	⬅	2 CREATE EXTERNAL DATA SOURCE
		3 CREATE EXTERNAL FILE FORMAT
		4 CREATE EXTERNAL TABLE

NEW QUESTION 114

DRAG DROP

You need to design a Microsoft Azure solution to analyze text from a Twitter data stream. The solution must identify a sentiment score of positive, negative, or neutral for the tweets.

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
Use the Bing Custom Search API to build a search.		
Call an anomaly detection function from the Stream Analytics job.		
Build an application that uses the Emotion API in Microsoft Cognitive Services.	➤	⬆
Create a user-defined function to call the sentiment analysis scoring web service.	⬅	⬇
Create an Azure Stream Analytics job.		
Deploy an Azure Machine Learning sentiment analytics model from the Cortana Intelligence Gallery as a web service.		

Answer:

Explanation:

Actions		Answer Area
Use the Bing Custom Search API to build a search.		
Call an anomaly detection function from the Stream Analytics job.		
Build an application that uses the Emotion API in Microsoft Cognitive Services.	➤	1 Deploy an Azure Machine Learning sentiment analytics model from the Cortana Intelligence Gallery as a web service.
	⬅	2 Create a user-defined function to call the sentiment analysis scoring web service.
		3 Create an Azure Stream Analytics job.

NEW QUESTION 117

.....

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>