



DP-200^{Q&As}

Implementing an Azure Data Solution

Pass Microsoft DP-200 Exam with 100% Guarantee

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

<https://www.passapply.com/dp-200.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

DRAG DROP

You have data stored in thousands of CSV files in Azure Data Lake Storage Gen2. Each file has a header row followed by a property formatted carriage return (/r) and line feed (/n).

You are implementing a pattern that batch loads the files daily into an enterprise data warehouse in Azure Synapse Analytics by using PolyBase.

You need to skip the header row when you import the files into the data warehouse.

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.

Select and Place:

Actions

Create an external data source that uses the abfs location.

Create an external file format and set the First_Row option.

Create an external data source that uses the Hadoop location.

Create a database scoped credential that uses an OAuth2 token and a key.

Use CREATE EXTERNAL TABLE AS SELECT (CETAS) and create a view that removes the empty row.

Use CREATE EXTERNAL TABLE AS SELECT (CETAS) and configure the reject options to specify reject values or percentages.

Answer Area

Three empty rectangular boxes with red borders, arranged vertically, intended for the user to place the selected actions in the correct order.

Correct Answer:



Actions

Create an external data source that uses the abfs location.

Create an external file format and set the First_Row option.

Create an external data source that uses the Hadoop location.

Create a database scoped credential that uses an OAuth2 token and a key.

Use CREATE EXTERNAL TABLE AS SELECT (CETAS) and create a view that removes the empty row.

Use CREATE EXTERNAL TABLE AS SELECT (CETAS) and configure the reject options to specify reject values or percentages.

Answer Area

Create an external file format and set the First_Row option.

Create an external data source that uses the abfs location.

Use CREATE EXTERNAL TABLE AS SELECT (CETAS) and create a view that removes the empty row.

Step 1: Create an external data source and set the First_Row option.

Creates an External File Format object defining external data stored in Hadoop, Azure Blob Storage, or Azure Data Lake Store. Creating an external file format is a prerequisite for creating an External Table.

FIRST_ROW = First_row_int

Specifies the row number that is read first in all files during a PolyBase load. This parameter can take values 1-15. If the value is set to two, the first row in every file (header row) is skipped when the data is loaded. Rows are skipped based on

the existence of row terminators (/r/n, /r, /n).

Step 2: Create an external data source that uses the abfs location

The hadoop-azure module provides support for the Azure Data Lake Storage Gen2 storage layer through the "abfs" connector

Step 3: Use CREATE EXTERNAL TABLE AS SELECT (CETAS) and create a view that removes the empty row.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-external-file-format-transact-sql>

<https://hadoop.apache.org/docs/r3.2.0/hadoop-azure/abfs.html>

QUESTION 2

DRAG DROP



You have an Azure Data Lake Storage Gen2 account that contains JSON files for customers. The files contain two attributes named FirstName and LastName.

You need to copy the data from the JSON files to an Azure SQL data Warehouse table by using Azure Databricks. A new column must be created that concatenates the FirstName and LastName values.

You create the following components:

1.
A destination table in SQL Data Warehouse
2.
An Azure Blob storage container
3.
A service principal

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

Select and Place:

Actions

- Write the results to Data Lake Storage.
- Drop the data frame.
- Perform transformations on the data frame.
- Mount the Data Lake Storage onto DBFS.
- Perform transformations on the file.
- Read the file into a data frame.
- Specify a temporary folder to stage the data.
- Write the results to a table in SQL Data Warehouse.

Answer Area

Five empty red-bordered boxes for placing the actions in sequence.

Correct Answer:



Actions

Write the results to Data Lake Storage.
Mount the Data Lake Storage onto DBFS.
Perform transformations on the file.

Answer Area

Read the file into a data frame.
Perform transformations on the data frame.
Specify a temporary folder to stage the data.
Write the results to a table in SQL Data Warehouse.
Drop the data frame.

Step 1: Read the file into a data frame.

You can load the json files as a data frame in Azure Databricks.

Step 2: Perform transformations on the data frame.

Step 3: Specify a temporary folder to stage the data

Specify a temporary folder to use while moving data between Azure Databricks and Azure SQL Data Warehouse.

Step 4: Write the results to a table in SQL Data Warehouse

You upload the transformed data frame into Azure SQL Data Warehouse. You use the Azure SQL Data Warehouse connector for Azure Databricks to directly upload a dataframe as a table in a SQL data warehouse.

Step 5: Drop the data frame

Clean up resources. You can terminate the cluster. From the Azure Databricks workspace, select Clusters on the left. For the cluster to terminate, under Actions, point to the ellipsis (...) and select the Terminate icon.

References:

<https://docs.microsoft.com/en-us/azure/azure-databricks/databricks-extract-load-sql-data-warehouse>

QUESTION 3

You develop data engineering solutions for a company.

You must integrate the company's on-premises Microsoft SQL Server data with Microsoft Azure SQL Database. Data must be transformed incrementally.



You need to implement the data integration solution.

Which tool should you use to configure a pipeline to copy data?

- A. Use the Copy Data tool with Blob storage linked service as the source
- B. Use Azure PowerShell with SQL Server linked service as a source
- C. Use Azure Data Factory UI with Blob storage linked service as a source
- D. Use the .NET Data Factory API with Blob storage linked service as the source

Correct Answer: C

The Integration Runtime is a customer managed data integration infrastructure used by Azure Data Factory to provide data integration capabilities across different network environments. A linked service defines the information needed for Azure Data Factory to connect to a data resource. We have three resources in this scenario for which linked services are needed: On-premises SQL Server Azure Blob Storage Azure SQL database

Note: Azure Data Factory is a fully managed cloud-based data integration service that orchestrates and automates the movement and transformation of data. The key concept in the ADF model is pipeline. A pipeline is a logical grouping of Activities, each of which defines the actions to perform on the data contained in Datasets. Linked services are used to define the information needed for Data Factory to connect to the data resources.

References: <https://docs.microsoft.com/en-us/azure/machine-learning/team-data-science-process/move-sql-azure-adf>

QUESTION 4

DRAG DROP

You have an Azure Active Directory (Azure AD) tenant that contains a security group named Group1. You have an Azure Synapse Analytics dedicated SQL pool named dw1 that contains a schema named schema1.

You need to grant Group1 read-only permissions to all the tables and views in schema1. The solution must use the principle of least privilege.

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.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Select and Place:



Actions

Answer Area

- Assign Role1 to the Group1 database user
- Create a database role named Role1 and grant Role1 SELECT permissions to schema1
- Create a database role named Role1 and grant Role1 SELECT permissions to dw1
- Create a database user in dw1 that represents Group1 and uses the FROM EXTERNAL PROVIDER clause
- Assign the Azure role-based access control (Azure RBAC) Reader role for dw1 to Group1




Correct Answer:

Actions

Answer Area

-
-
- Create a database role named Role1 and grant Role1 SELECT permissions to dw1
- Create a database user in dw1 that represents Group1 and uses the FROM EXTERNAL PROVIDER clause
-

- Create a database role named Role1 and grant Role1 SELECT permissions to schema1
- Assign Role1 to the Group1 database user
- Assign the Azure role-based access control (Azure RBAC) Reader role for dw1 to Group1



Step 1: Create a database role named Role1 and grant Role1 SELECT permissions to schema

You need to grant Group1 read-only permissions to all the tables and views in schema1.

Place one or more database users into a database role and then assign permissions to the database role.

Step 2: Assign Rol1 to the Group database user



Step 3: Assign the Azure role-based access control (Azure RBAC) Reader role for dw1 to Group1

Reference:

<https://docs.microsoft.com/en-us/azure/data-share/how-to-share-from-sql>

QUESTION 5

You use Azure Stream Analytics to receive Twitter data from Azure Event Hubs and to output the data to an Azure Blob storage account.

You need to output the count of tweets during the last five minutes every five minutes. Each tweet must only be counted once.

Which windowing function should you use?

- A. a five-minute Session window
- B. a five-minute Sliding window
- C. a five-minute Tumbling window
- D. a five-minute Hopping window that has one-minute hop

Correct Answer: C

Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The key differentiators of a Tumbling window are that they repeat, do not overlap, and an event cannot belong to more than one tumbling window.



Tell me the count of tweets per time zone every 10 seconds



```
SELECT TimeZone, COUNT(*) AS Count  
FROM TwitterStream TIMESTAMP BY CreatedAt  
GROUP BY TimeZone, TumblingWindow(second,10)
```

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

[Latest DP-200 Dumps](#)

[DP-200 Exam Questions](#)

[DP-200 Braindumps](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

- 100% Guaranteed Success
- 100% Money Back Guarantee
- 365 Days Free Update
- Instant Download After Purchase
- 24x7 Customer Support
- Average 99.9% Success Rate
- More than 800,000 Satisfied Customers Worldwide
- Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.passapply.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket:



 <p>One Year Free Update Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p>Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p>Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.
All trademarks are the property of their respective owners.
Copyright © passapply, All Rights Reserved.