



Vendor: Microsoft

Exam Code: 70-776

Exam Name: Performing Big Data Engineering with
Microsoft Cloud Services

Version: DEMO

QUESTION 1

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 folders. Configure the queries to use uniform partition keys.
- C. Set the partition key for the inputs, queries, and outputs to use the same partition folders. Configure the queries to use different partition keys.
- D. Define the number of input partitions to equal the number of output partitions.

Answer: A

Explanation:

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

QUESTION 2

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

QUESTION 3

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:

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

QUESTION 4

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:

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

QUESTION 5

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

QUESTION 6

You need to use the Cognition.Vision.FaceDetector() function in U-SQL to analyze images.

Which attribute can you detect by using the function?

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

Answer: A

QUESTION 7

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

Thank You for Trying Our Product

PassLeader Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.passleader.com/all-products.html>



10% Discount Coupon Code: STNAR2014