# Question 1

You have two Azure Blob Storage accounts named account1 and account2?

You plan to create an Azure Data Factory pipeline that will use scheduled intervals to replicate newly created or modified blobs from account1 to account?

You need to recommend a solution to implement the pipeline. The solution must meet the following requirements:

* Ensure that the pipeline only copies blobs that were created of modified since the most recent replication event.

* Minimize the effort to create the pipeline.

What should you recommend?

## Options:

A- Create a pipeline that contains a flowlet.

B- Create a pipeline that contains a Data Flow activity.

C- Run the Copy Data tool and select Metadata-driven copy task.

D- Run the Copy Data tool and select Built-in copy task.

**Answer:**

A

# Question 2

Question Type: **MultipleChoice**

Note: The 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 an Azure Data Lake Storage account that contains a staging zone.

You need to design a dairy process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.

Solution: You use an Azure Data Factory schedule trigger to execute a pipeline that executes a mapping data low. and then inserts the data into the data warehouse.

Does this meet the goal?

**Options:**

**A-** Yes

**B-** No

## Answer:

B

# Question 3

**Question Type: MultipleChoice**

You have an Azure data factory named ADM that contains a pipeline named Pipelwe1

Pipeline! must execute every 30 minutes with a 15-minute offset.

Vou need to create a trigger for Pipehne1. The trigger must meet the following requirements:

* Backfill data from the beginning of the day to the current time.

* If Pipeline1 fairs, ensure that the pipeline can re-execute within the same 30-mmute period.

* Ensure that only one concurrent pipeline execution can occur.

* Minimize de4velopment and configuration effort

Which type of trigger should you create?

## Options:

**A-** schedule

**B-** event-based

**C-** manual

**D-** tumbling window

## Answer:

A

# Question 4

**Question Type:** **MultipleChoice**

You have an Azure subscription that contains an Azure Synapse Analytics workspace named ws1 and an Azure Cosmos D6 database account named Cosmos1 Costmos1 contains a container named container 1 and ws1 contains a serverless1 SQL pool.

you need to ensure that you can Query the data in container by using the serverless1 SQL pool.

Which three actions should you perform? Each correct answer presents part of the solution

NOTE: Each correct selection is worth one point.

## Options:

**A-** Enable Azure Synapse Link for Cosmos1

**B-** Disable the analytical store for container1.

**C-** In ws1. create a linked service that references Cosmos1

**D-** Enable the analytical store for container1

**E-** Disable indexing for container1

## Answer:

A, C, D

# Question 5

**Question Type:** **MultipleChoice**

You have an Azure Synapse Analytics dedicated SQL pod.

You need to create a pipeline that will execute a stored procedure in the dedicated SQL pool and use the returned result set as the input (or a downstream activity. The solution must minimize development effort.

Which Type of activity should you use in the pipeline?

## Options:

**A-** Notebook

**B-** U-SQL

**C-** Script

**D-** Stored Procedure

## Answer:

D

# Question 6

You are building a data flow in Azure Data Factory that upserts data into a table in an Azure Synapse Analytics dedicated SQL pool.

You need to add a transformation to the data flow. The transformation must specify logic indicating when a row from the input data must be upserted into the sink.

Which type of transformation should you add to the data flow?

## Options:

**A-** join

**B-** select

**C-** surrogate key

**D-** alter row

## Answer:

D

## Explanation:

The alter row transformation allows you to specify insert, update, delete, and upsert policies on rows based on expressions. You can use the alter row transformation to perform upserts on a sink table by matching on a key column and setting the appropriate row policy
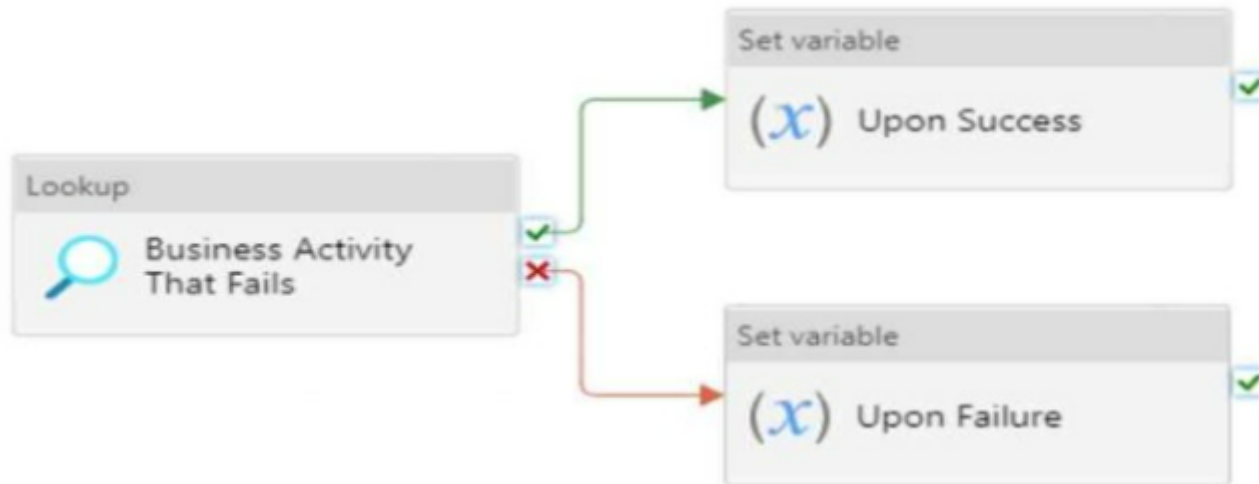
# Question 7

You have the Azure Synapse Analytics pipeline shown in the following exhibit.



You need to add a set variable activity to the pipeline to ensure that after the pipeline's completion, the status of the pipeline is always successful.

What should you configure for the set variable activity?

## Options:

**A-** a success dependency on the Business Activity That Fails activity

**B-** a failure dependency on the Upon Failure activity

**C-** a skipped dependency on the Upon Success activity

**D-** a skipped dependency on the Upon Failure activity

## Answer:

A

## Explanation:

A failure dependency means that the activity will run only if the previous activity fails.In this case, setting a failure dependency on the Upon Failure activity will ensure that the set variable activity will run after the pipeline fails and set the status of the pipeline to successful.

# Question 8

**Question Type:** **MultipleChoice**

You have an Azure subscription that contains a Microsoft Purview account named MP1, an Azure data factory named DF1, and a storage account named storage. MP1 is configured

10 scan storage1. DF1 is connected to MP1 and contains 3 dataset named DS1. DS1 references 2 file in storage.

In DF1, you plan to create a pipeline that will process data from DS1.

You need to review the schema and lineage information in MP1 for the data referenced by DS1.

Which two features can you use to locate the information? Each correct answer presents a complete solution.

NOTE: Each correct answer is worth one point.

## Options:

**A-** the Storage browser of storage1 in the Azure portal

**B-** the search bar in the Azure portal

**C-** the search bar in Azure Data Factory Studio

**D-** the search bar in the Microsoft Purview governance portal

## Answer:

C, D

## Explanation:

The search bar in the Microsoft Purview governance portal: This is a feature that allows you to search for assets in your data estate using keywords, filters, and facets.You can use the search bar to find the files in storage1 that are referenced by DS1, and then view their schema and lineage information in the asset details page12.

The search bar in Azure Data Factory Studio: This is a feature that allows you to search for datasets, linked services, pipelines, and other resources in your data factory. You can use the search bar to find DS1 in DF1, and then view its schema and lineage information in the dataset details page.You can also click on the Open in Purview button to open the corresponding asset in MP13.

The two features that can be used to locate the schema and lineage information for the data referenced by DS1 are the search bar in Azure Data Factory Studio and the search bar in the Microsoft Purview governance portal.

The search bar in Azure Data Factory Studio allows you to search for the dataset DS1 and view its properties and lineage. This can help you locate information about the source and destination data stores, as well as the transformations that were applied to the data.

The search bar in the Microsoft Purview governance portal allows you to search for the storage account and view its metadata, including schema and lineage information. This can help you understand the different data assets that are stored in the storage account and how they are related to each other.

The Storage browser of storage1 in the Azure portal may allow you to view the files that are stored in the storage account, but it does not provide lineage or schema information for those files. Similarly, the search bar in the Azure portal may allow you to search for resources in the Azure subscription, but it does not provide detailed information about the data assets themselves.

What is Azure Purview?

Use Azure Data Factory Studio

# Question 9

You are deploying a lake database by using an Azure Synapse database template.

You need to add additional tables to the database. The solution must use the same grouping method as the template tables.

'Which grouping method should you use?

## Options:

**A-** business area

**B-** size

**C-** facts and dimensions

**D-** partition style

## Answer:

A

**Explanation:**

Business area: This is how the Azure Synapse database templates group tables by default. Each template consists of one or more enterprise templates that contain tables grouped by business areas.For example, the Retail template has business areas such as Customer, Product, Sales, and Store123. Using the same grouping method as the template tables can help you maintain consistency and compatibility with the industry-specific data model.

https://techcommunity.microsoft.com/t5/azure-synapse-analytics-blog/database-templates-in-azure-synapse-analytics/ba-p/2929112