A.
```
DROP TABLE [Ext].[Items]
CREATE EXTERNAL TABLE [Ext].[Items]
( [ItemID] [int] NULL,
  [ItemName] nvarchar(50) NULL,
  [ItemType] nvarchar(20) NULL,
  [ItemDescription] nvarchar(250))
WITH
(
    LOCATION='/Items/',
      DATA_SOURCE = AzureDataLakeStore,
      FILE_FORMAT = PARQUET,
      REJECT_TYPE = VALUE,
      REJECT_VALUE = 0
);
```
B.
```
ALTER EXTERNAL TABLE [Ext].[Items]
ADD [ItemID] int;
```
C.
```
DROP EXTERNAL FILE FORMAT parquetfile1;
CREATE EXTERNAL FILE FORMAT parquetfile1
WITH (
    FORMAT_TYPE = PARQUET,
    DATA_COMPRESSION = 'org.apache.hadoop.io.compress.SnappyCodec'
);
```
D.
```
ALTER TABLE [Ext].[Items]
ADD [ItemID] int
```

A. Option A
B. Option B
C. Option C
D. Option D

Correct Answer: A
Explanation:
https://docs.microsoft.com/en-us/sql/t-sql/statements/create-external-table-transact-sql

QUESTION 83
You plan to create a dimension table in Azure Data Warehouse that will be less than 1 GB.

You need to create the table to meet the following requirements:

- Provide the fastest query time.
- Minimize data movement.

Which type of table should you use?

A. hash distributed
B. heap
C. replicated
D. round-robin

Correct Answer: D
Explanation:
Usually common dimension tables or tables that doesn't distribute evenly are good candidates for round-robin distributed table.

Note:
Dimension tables or other lookup tables in a schema can usually be stored as round-robin tables.

Usually these tables connect to more than one fact tables and optimizing for one join may not be the best idea.

Also usually dimension tables are smaller which can leave some distributions empty when hash distributed.

Round-robin by definition guarantees a uniform data distribution.

References:
https://blogs.msdn.microsoft.com/sqlcat/2015/08/11/choosing-hash-distributed-table-vs-round-robindistributed-table-in-azure-sql-dw-service/

QUESTION 84
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 scenario, 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 Storage account that contains 100 GB of files. The files contain text and numerical values. 75% of the rows contain description data that has an average length of 1.1 MB.

You plan to copy the data from the storage account to an Azure SQL data warehouse.

You need to prepare the files to ensure that the data copies quickly.

Solution: You copy the files to a table that has a columnstore index.

Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Explanation:
Instead modify the files to ensure that each row is less than 1 MB.

References:

https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data

QUESTION 85
DRAG DROP
You have a table named SalesFact in an Azure SQL data warehouse. SalesFact contains sales data from the past 36 months and has the following characteristics:

- Is partitioned by month
- Contains one billion rows
- Has clustered columnstore indexes

All the beginning of each month, you need to remove data SalesFact that is older than 36 months as quickly as possible.

Which three actions should you perform in sequence in a stored procedure? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

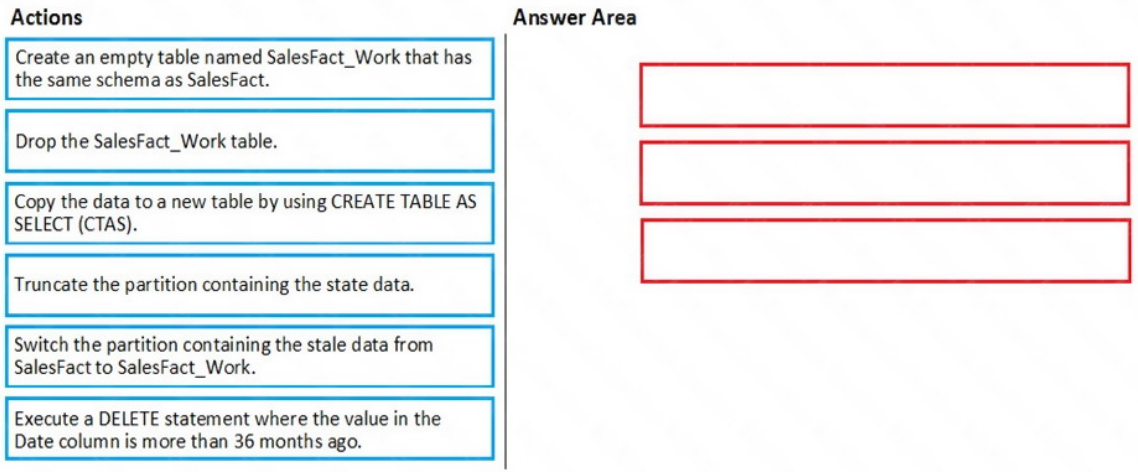

Correct Answer:

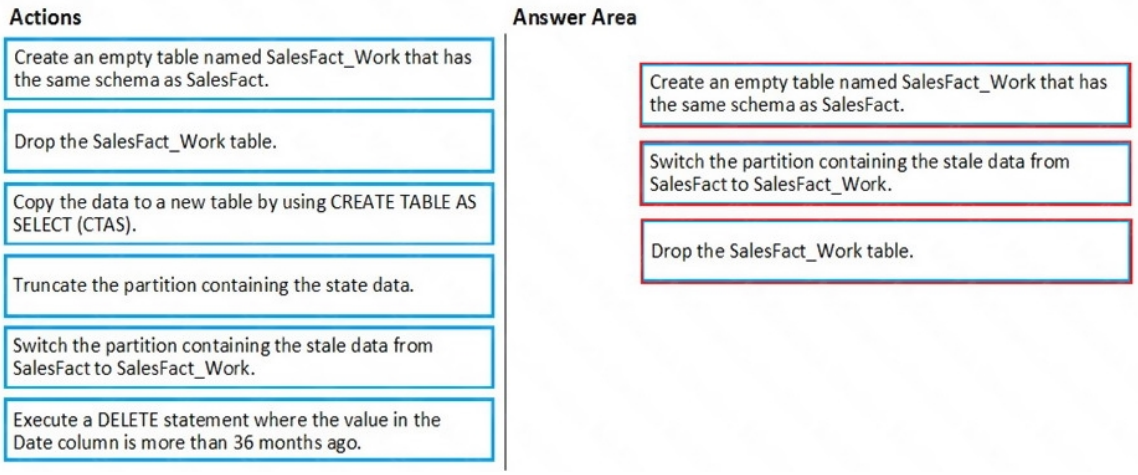

QUESTION 86
DRAG DROP
You plan to monitor an Azure data factory by using the Monitor & Manage app.

You need to identify the status and duration of activities that reference a table in a source database.

Which three actions should you perform in sequence? To answer, move the actions from the list of actions to the answer are and arrange them in the correct order.

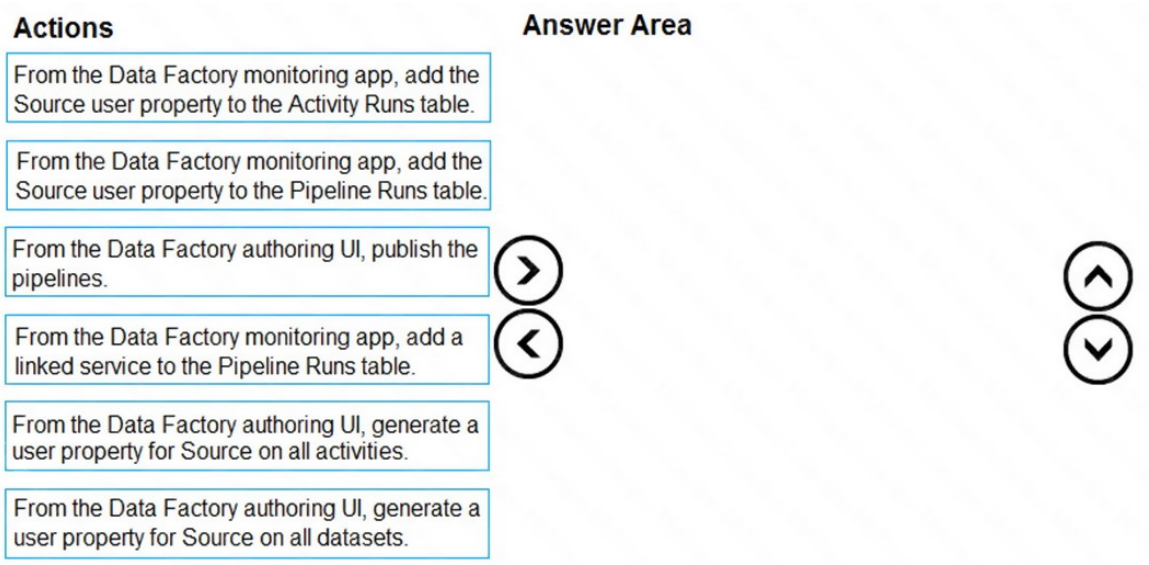

Correct Answer:

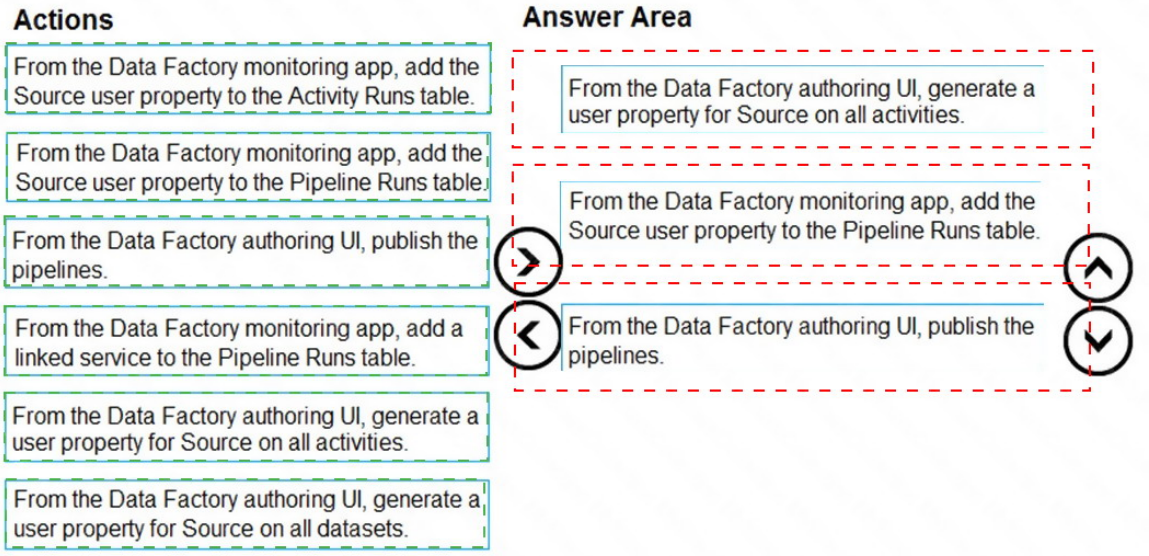

QUESTION 87
You have an Azure Storage account and an Azure SQL data warehouse by using Azure Data Factory. The solution must meet the following requirements:

- Ensure that the data remains in the UK South region at all times.
- Minimize administrative effort.

Which type of integration runtime should you use?

A. Azure integration runtime
B. Self-hosted integration runtime
C. Azure-SSIS integration runtime
D. None of the above

Correct Answer: A
Explanation:
https://docs.microsoft.com/en-us/azure/data-factory/concepts-integration-runtime

QUESTION 88
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 scenario, 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 Storage account that contains 100 GB of files. The files contain text and numerical values. 75% of the rows contain description data that has an average length of 1.1 MB.

You plan to copy the data from the storage account to an Azure SQL data warehouse.

You need to prepare the files to ensure that the data copies quickly.

Solution: You modify the files to ensure that each row is less than 1 MB.

Does this meet the goal?

A. Yes
B. No

Correct Answer: A
Explanation:
When exporting data into an ORC File Format, you might get Java out-of-memory errors when there are large text columns. To work around this limitation, export only a subset of the columns.

References:
https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data