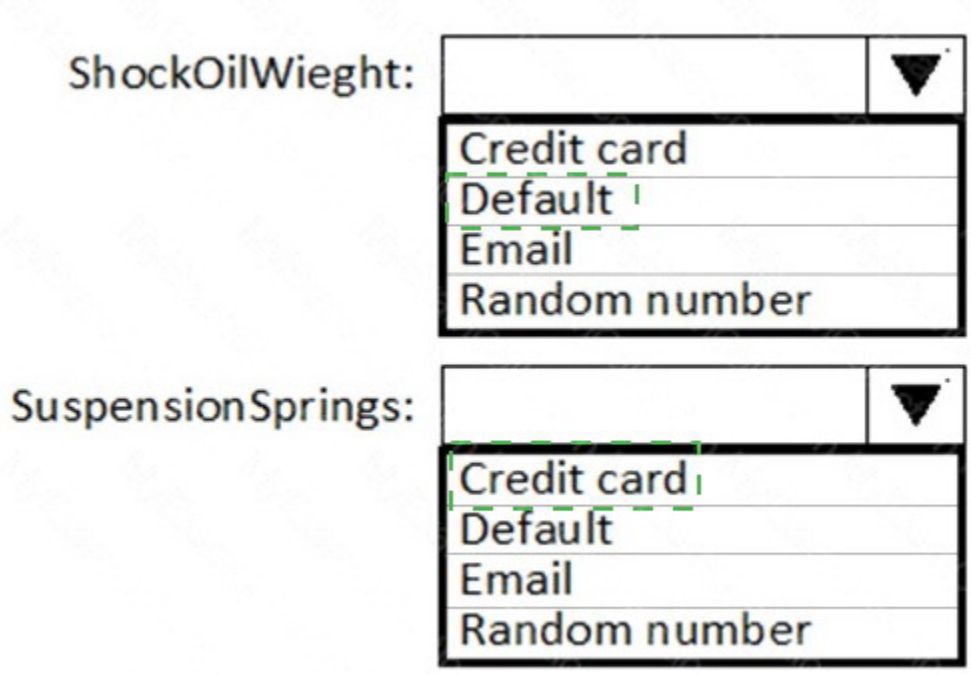

QUESTION 2
You are monitoring the Data Factory pipeline that runs from Cosmos DB to SQL Database for Race Central.

You discover that the job takes 45 minutes to run.

What should you do to improve the performance of the job?

A. Decrease parallelism for the copy activities.
B. Increase that data integration units.
C. Configure the copy activities to use staged copy.
D. Configure the copy activities to perform compression.

Correct Answer: B
Explanation:
Performance tuning tips and optimization features. In some cases, when you run a copy activity in Azure Data Factory, you see a "Performance tuning tips" message on top of the copy activity monitoring, as shown in the following example. The message tells you the bottleneck that was identified for the given copy run. It also guides you on what to change to boost copy throughput. The performance tuning tips currently provide suggestions like:

Use PolyBase when you copy data into Azure SQL Data Warehouse.

Increase Azure Cosmos DB Request Units or Azure SQL Database DTUs (Database Throughput Units) when the resource on the data store side is the bottleneck.

Remove the unnecessary staged copy.

References:
https://docs.microsoft.com/en-us/azure/data-factory/copy-activity-performance

QUESTION 3
HOTSPOT
You are building the data store solution for Mechanical Workflow.

How should you configure Table1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

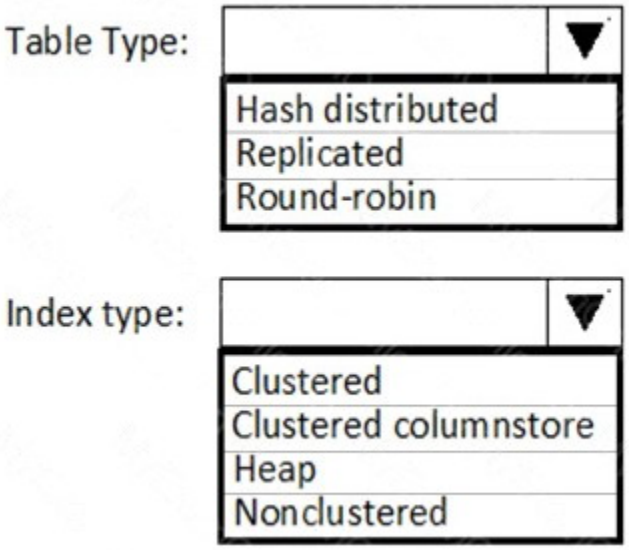

Correct Answer:

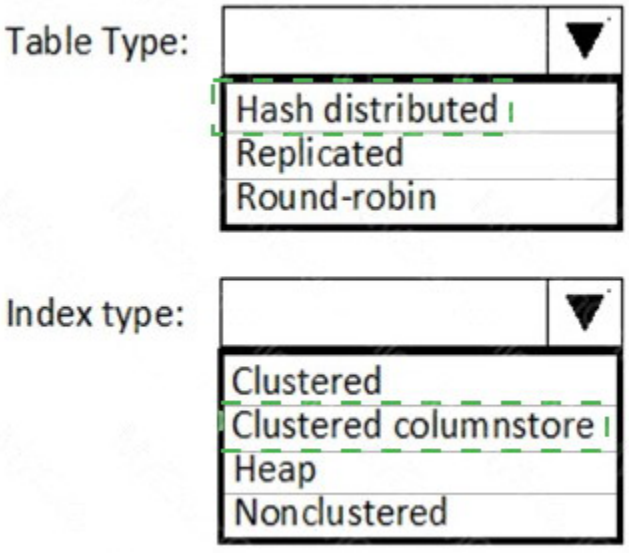

QUESTION 4

On which data store you configure TDE to meet the technical requirements?

A. Cosmos DB
B. SQL Data Warehouse
C. SQL Database
D. None of the above

Correct Answer: B
Explanation:
Scenario: Transparent data encryption (TDE) must be enabled on all data stores, whenever possible.
The datacentre for Mechanical Workflow must be moved to Azure SQL data Warehouse.

QUESTION 5
Which two metrics should you use to identify the appropriate RU/s for the telemetry data? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Number of requests
B. Number of requests exceeded capacity
C. End to end observed read latency at the 99th percentile
D. Session consistency
E. Data + Index storage consumed
F. Avg Troughput/s

Correct Answer: AE
Explanation:
Scenario: The telemetry data must be monitored for performance issues. You must adjust the Cosmos DB Request Units per second (RU/s) to maintain a performance SLA while minimizing the cost of the Ru/s.

With Azure Cosmos DB, you pay for the throughput you provision and the storage you consume on an hourly basis.

While you estimate the number of RUs per second to provision, consider the following factors:

Item size: As the size of an item increases, the number of RUs consumed to read or write the item also increases.

QUESTION 6
HOTSPOT
You need to build a solution to collect the telemetry data for Race Control.

What should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

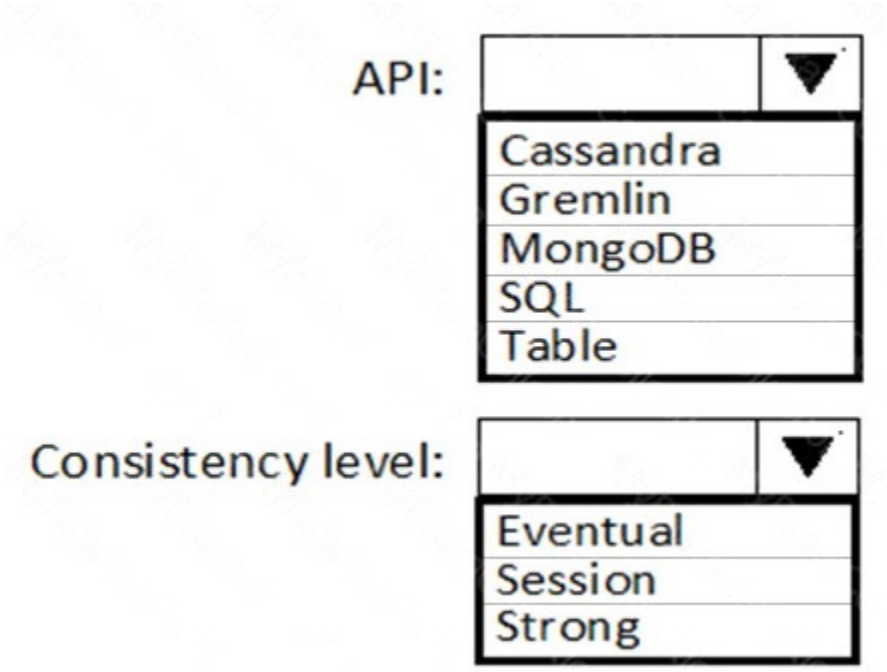

Correct Answer:

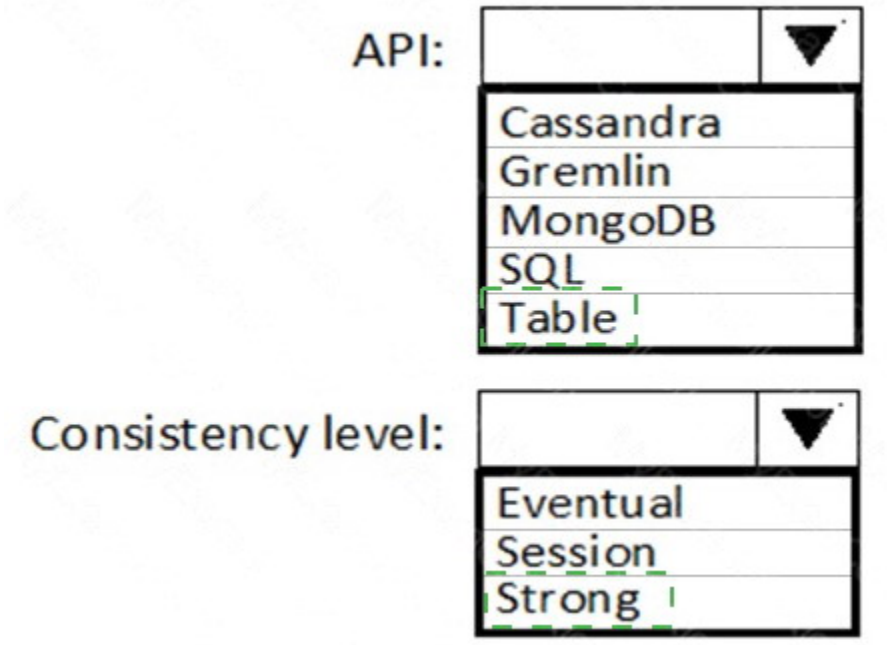

QUESTION 7
What should you implement to optimize SQL Database for Race Central to meet the technical requirements?

A. the sp_update stored procedure
B. automatic tuning
C. Query Store
D. the dbcc checkdb command

Correct Answer: A
Explanation:
Scenario: The query performance of Race Central must be stable, and the administrative time it takes to perform optimizations must be minimized.

sp_update updates query optimization statistics on a table or indexed view. By default, the query optimizer already updates statistics as necessary to improve the query plan; in some cases you can improve query performance by using UPDATE STATISTICS or the stored procedure sp_updatestats to update statistics more frequently than the default updates.

QUESTION 8
What should you include in the Data Factory pipeline for Race Central?

A. a copy activity that uses a stored procedure as a source
B. a copy activity that contains schema mappings
C. a delete activity that has logging enabled
D. a filter activity that has a condition

Correct Answer: B
Explanation:
Scenario:
An Azure Data Factory pipeline must be used to move data from Cosmos DB to SQL Database for Race Central. If the data load takes longer than 20 minutes, configuration changes must be made to Data Factory.

The telemetry data is sent to a MongoDB database. A custom application then moves the data to databases in SQL Server 2017. The telemetry data in MongoDB has more than 500 attributes. The application changes the attribute names when the data is moved to SQL Server 2017.

You can copy data to or from Azure Cosmos DB (SQL API) by using Azure Data Factory pipeline.

Column mapping applies when copying data from source to sink. By default, copy activity map source data to sink by column names. You can specify explicit mapping to customize the column mapping based on your need. More specifically, copy activity:

Read the data from source and determine the source schema

Use default column mapping to map columns by name, or apply explicit column mapping if specified.
Write the data to sink
Write the data to sink