DP-201 Dumps

Designing an Azure Data Solution

https://www.certleader.com/DP-201-dumps.html

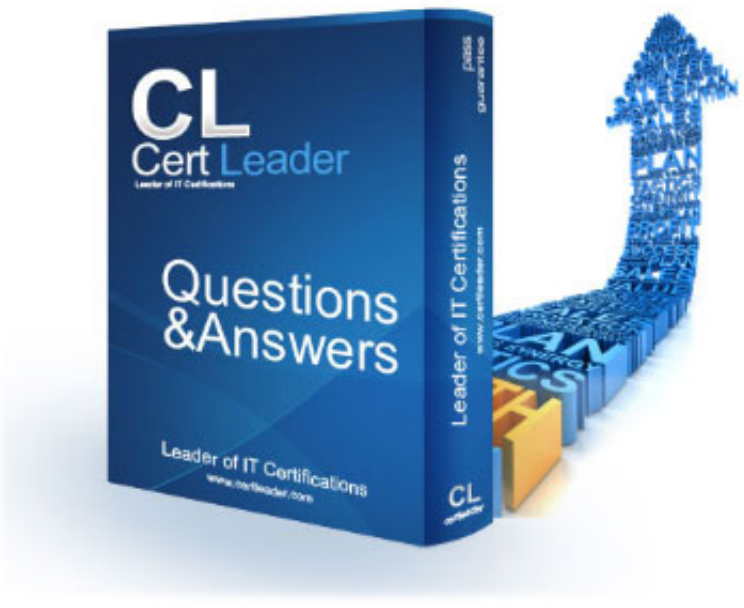

NEW QUESTION 1
- (Exam Topic 1)
You need to design the Planning Assistance database.
For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Statement	Yes	No
Including a clustered columnstore index in the design will benefit performance.	○	○
Including a nonclustered columnstore index in the design will benefit performance.	○	○
Including an index on the License Plate column will benefit performance.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: No
Data used for Planning Assistance must be stored in a sharded Azure SQL Database. Box 2: Yes
Box 3: Yes
Planning Assistance database will include reports tracking the travel of a single vehicle

NEW QUESTION 2
- (Exam Topic 1)
You need to ensure that emergency road response vehicles are dispatched automatically.
How should you design the processing system? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

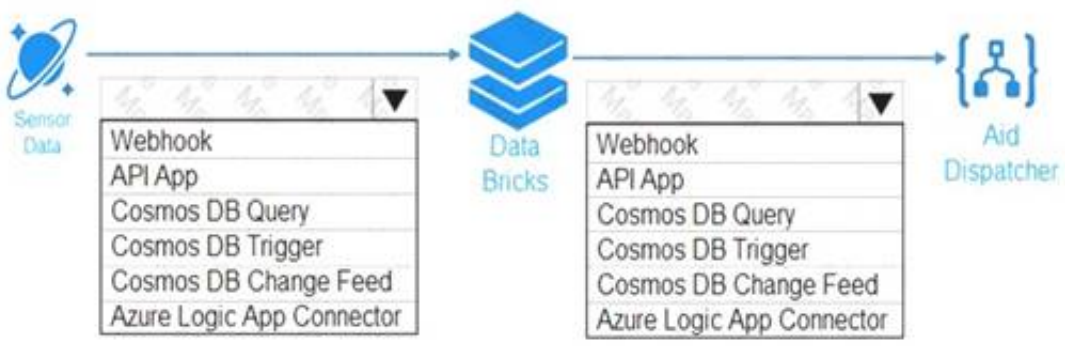

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box1: API App

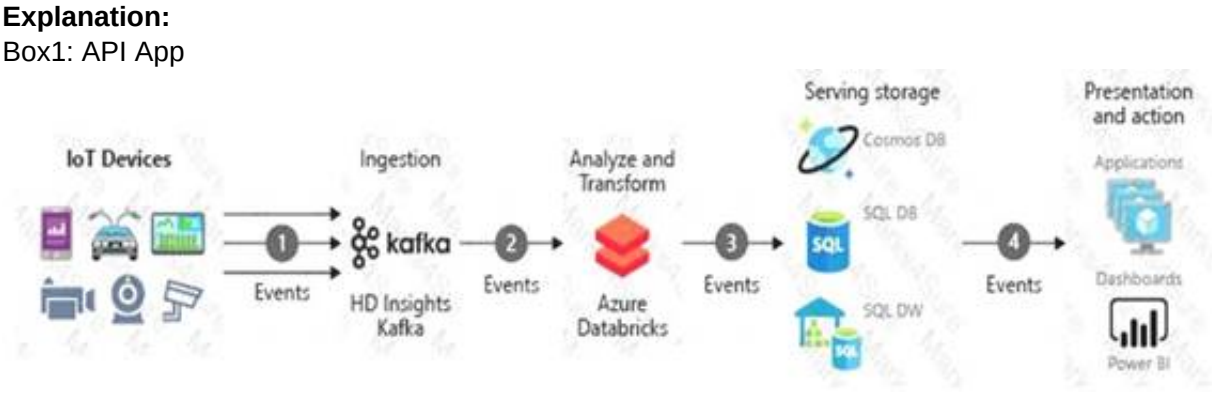

- Events generated from the IoT data sources are sent to the stream ingestion layer through Azure HDInsight Kafka as a stream of messages. HDInsight Kafka stores streams of data in topics for a configurable of time.
- Kafka consumer, Azure Databricks, picks up the message in real time from the Kafka topic, to process the data based on the business logic and can then send to Serving layer for storage.
- Downstream storage services, like Azure Cosmos DB, Azure SQL Data warehouse, or Azure SQL DB, will then be a data source for presentation and action layer.
- Business analysts can use Microsoft Power BI to analyze warehoused data. Other applications can be built upon the serving layer as well. For example, we can expose APIs based on the service layer data for third party uses.

Box 2: Cosmos DB Change Feed
Change feed support in Azure Cosmos DB works by listening to an Azure Cosmos DB container for any changes. It then outputs the sorted list of documents that were changed in the order in which they were modified.
The change feed in Azure Cosmos DB enables you to build efficient and scalable solutions for each of these patterns, as shown in the following image:

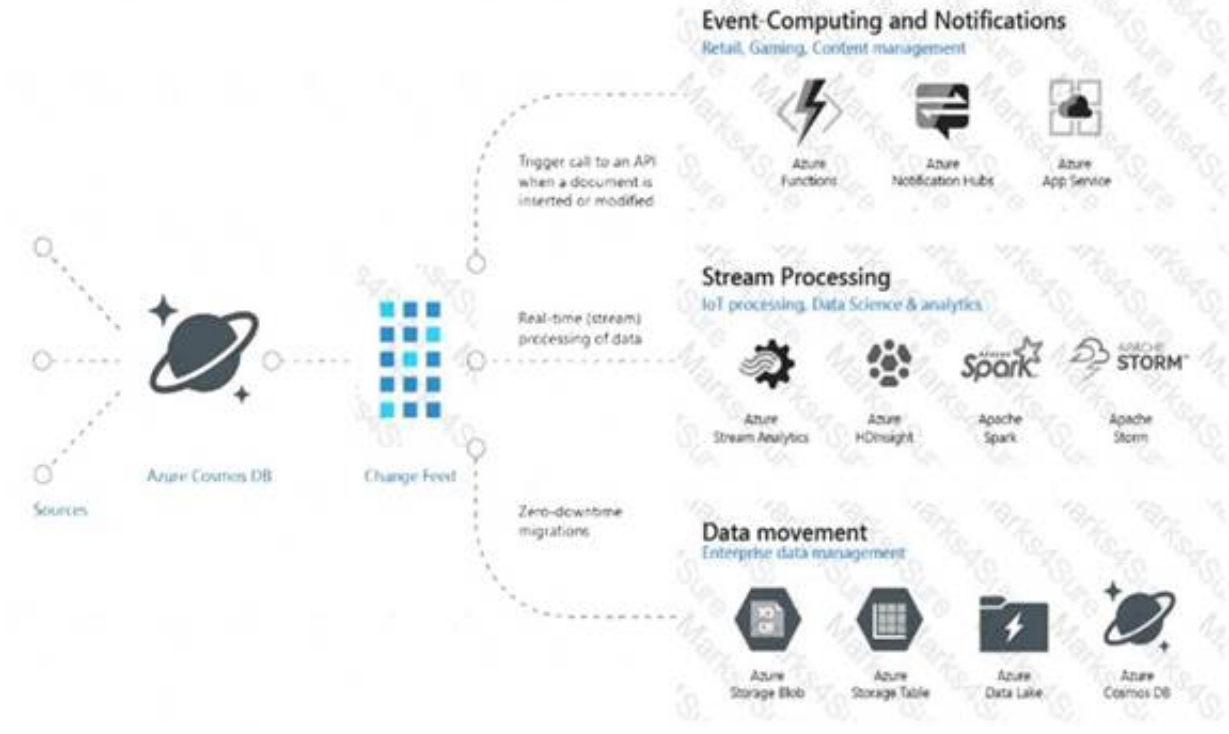

References:
https://docs.microsoft.com/bs-cyrl-ba/azure/architecture/example-scenario/data/realtime-analytics-vehicle-iot?vi

NEW QUESTION 3
- (Exam Topic 1)
HOTSPOT
You need to ensure that security policies for the unauthorized detection system are met. What should you recommend? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

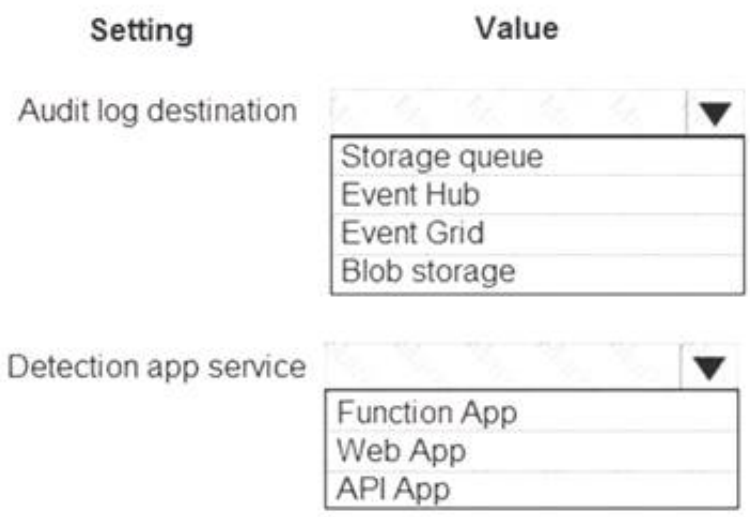

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Blob storage
Configure blob storage for audit logs.
Scenario: Unauthorized usage of the Planning Assistance data must be detected as quickly as possible. Unauthorized usage is determined by looking for an unusual pattern of usage.
Data used for Planning Assistance must be stored in a sharded Azure SQL Database. Box 2: Web Apps
SQL Advanced Threat Protection (ATP) is to be used.
One of Azure's most popular service is App Service which enables customers to build and host web applications in the programming language of their choice without managing infrastructure. App Service offers auto-scaling and high availability, supports both Windows and Linux. It also supports automated deployments from GitHub, Visual Studio Team Services or any Git repository. At RSA, we announced that Azure Security Center leverages the scale of the cloud to identify attacks targeting App Service applications.
References:
https://azure.microsoft.com/sv-se/blog/azure-security-center-can-identify-attacks-targeting-azure-app-service-ap

NEW QUESTION 4
- (Exam Topic 1)
You need to design the SensorData collection.
What should you recommend? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

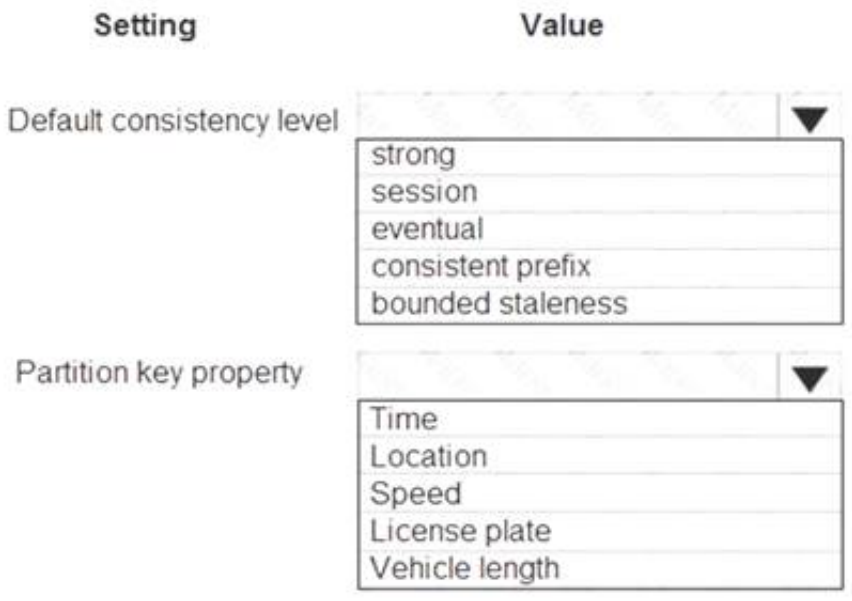

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Eventual
Traffic data insertion rate must be maximized.
Sensor data must be stored in a Cosmos DB named treydata in a collection named SensorData
With Azure Cosmos DB, developers can choose from five well-defined consistency models on the consistency spectrum. From strongest to more relaxed, the models include strong, bounded staleness, session, consistent prefix, and eventual consistency.
Box 2: License plate
This solution reports on all data related to a specific vehicle license plate. The report must use data from the SensorData collection.
References:
https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels

NEW QUESTION 5
- (Exam Topic 1)
STION NO: 5 HOTSPOT
You need to design the authentication and authorization methods for sensors.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

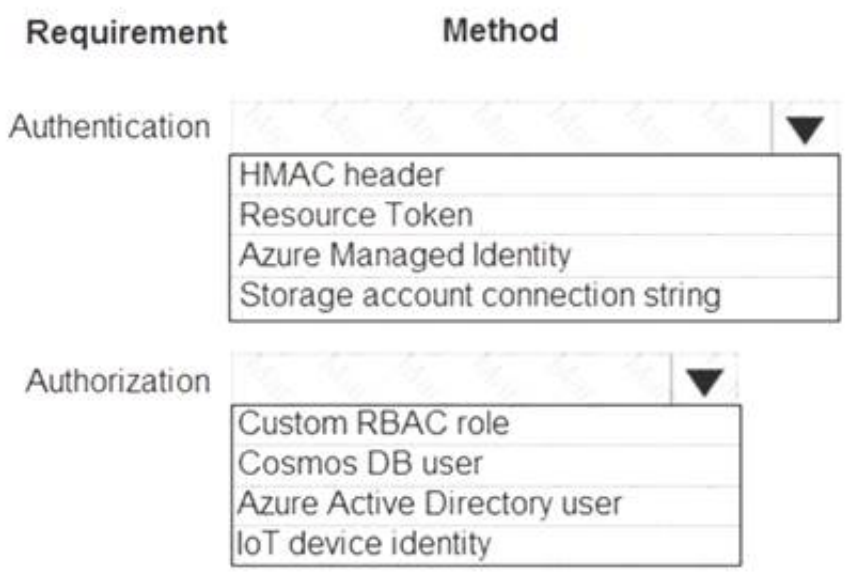

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Sensor data must be stored in a Cosmos DB named treydata in a collection named SensorData Sensors must have permission only to add items to the SensorData collection
Box 1: Resource Token
Resource tokens provide access to the application resources within a Cosmos DB database.
Enable clients to read, write, and delete resources in the Cosmos DB account according to the permissions they've been granted.
Box 2: Cosmos DB user
You can use a resource token (by creating Cosmos DB users and permissions) when you want to provide access to resources in your Cosmos DB account to a client that cannot be trusted with the master key.
References:
https://docs.microsoft.com/en-us/azure/cosmos-db/secure-access-to-data

NEW QUESTION 6
- (Exam Topic 2)
You need to design the solution for analyzing customer data. What should you recommend?

A. Azure Databricks
B. Azure Data Lake Storage
C. Azure SQL Data Warehouse
D. Azure Cognitive Services
E. Azure Batch

Answer: A

Explanation:
Customer data must be analyzed using managed Spark clusters. You create spark clusters through Azure Databricks. References:
https://docs.microsoft.com/en-us/azure/azure-databricks/quickstart-create-databricks-workspace-portal

NEW QUESTION 7
- (Exam Topic 2)
You need to design the image processing and storage solutions.
What should you recommend? To answer, select the appropriate configuration in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

Solution component	Tool
Image processing	Azure HDInsight Azure Databricks Azure Batch Azure Cognitive Services
data storage for tagging data	Azure Blob Storage Azure Table Storage Azure Cosmos DB Azure SQL Database

A. Mastered
B. Not Mastered

Answer: A

Explanation:
References:
https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/batch-processing https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tier-hyperscale

NEW QUESTION 8
- (Exam Topic 3)
A company stores sensitive information about customers and employees in Azure SQL Database. You need to ensure that the sensitive data remains encrypted in transit and at rest.
What should you recommend?

A. Transparent Data Encryption
B. Always Encrypted with secure enclaves
C. Azure Disk Encryption
D. SQL Server AlwaysOn

Answer: B

Explanation:
References:
https://cloudblogs.microsoft.com/sqlserver/2018/12/17/confidential-computing-using-always-encrypted-withsec

NEW QUESTION 9
- (Exam Topic 3)
You need to recommend a backup strategy for CONT_SQL1 and CONT_SQL2. What should you recommend?

A. Use AzCopy and store the data in Azure.
B. Configure Azure SQL Database long-term retention for all databases.
C. Configure Accelerated Database Recovery.
D. Use DWLoader.

Answer: B

Explanation:
Scenario: The database backups have regulatory purposes and must be retained for seven years.

NEW QUESTION 10
- (Exam Topic 4)
You design data engineering solutions for a company.
You must integrate on-premises SQL Server data into an Azure solution that performs Extract-Transform-Load (ETL) operations have the following requirements:
- Develop a pipeline that can integrate data and run notebooks.
- Develop notebooks to transform the data.
-

Load the data into a massively parallel processing database for later analysis. You need to recommend a solution.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 10
- (Exam Topic 4)
A company has many applications. Each application is supported by separate on-premises databases. You must migrate the databases to Azure SQL Database.
You have the following requirements: Organize databases into groups based on database usage.
Define the maximum resource limit available for each group of databases.
You need to recommend technologies to scale the databases to support expected increases in demand. What should you recommend?

A. Read scale-out
B. Managed instances
C. Elastic pools
D. Database sharding

Answer: C

Explanation:
SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single Azure SQL Database server and share a set number of resources at a set price.
You can configure resources for the pool based either on the DTU-based purchasing model or the vCorebased purchasing model.

NEW QUESTION 12
- (Exam Topic 4)
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 section, you will NOT be able to return to it. As a result, these
questions will not appear in the review screen.
You are designing an HDInsight/Hadoop cluster solution that uses Azure Data Lake Gen1 Storage. The solution requires POSIX permissions and enables diagnostics logging for auditing.
You need to recommend solutions that optimize storage.
Proposed Solution: Ensure that files stored are larger than 250MB. Does the solution meet the goal?

A. Yes
B. No

Answer: A

Explanation:
Depending on what services and workloads are using the data, a good size to consider for files is 256 MB or greater. If the file sizes cannot be batched when landing in Data Lake Storage Gen1, you can have a separate compaction job that combines these files into larger ones.
Note: POSIX permissions and auditing in Data Lake Storage Gen1 comes with an overhead that becomes apparent when working with numerous small files. As a best practice, you must batch your data into larger files versus writing thousands or millions of small files to Data Lake Storage Gen1. Avoiding small file sizes can have multiple benefits, such as:
Lowering the authentication checks across multiple files Reduced open file connections
Faster copying/replication
Fewer files to process when updating Data Lake Storage Gen1 POSIX permissions References:
https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-best-practices

NEW QUESTION 15
- (Exam Topic 4)
You plan to deploy an Azure SQL Database instance to support an application. You plan to use the DTUbased purchasing model.
Backups of the database must be available for 30 days and point-in-time restoration must be possible. You need to recommend a backup and recovery policy.
What are two possible ways to achieve the goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Use the Premium tier and the default backup retention policy.
B. Use the Basic tier and the default backup retention policy.
C. Use the Standard tier and the default backup retention policy.
D. Use the Standard tier and configure a long-term backup retention policy.
E. Use the Premium tier and configure a long-term backup retention policy.

Answer: DE

Explanation:
The default retention period for a database created using the DTU-based purchasing model depends on the service tier:
- Basic service tier is 1 week.
- Standard service tier is 5 weeks.
- Premium service tier is 5 weeks.

References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-long-term-retention

NEW QUESTION 18
- (Exam Topic 4)
You have an on-premises MySQL database that is 800 GB in size.
You need to migrate a MySQL database to Azure Database for MySQL. You must minimize service interruption to live sites or applications that use the database.
What should you recommend?

A. Azure Database Migration Service
B. Dump and restore
C. Import and export
D. MySQL Workbench

Answer: A

Explanation:
You can perform MySQL migrations to Azure Database for MySQL with minimal downtime by using the newly introduced continuous sync capability for the Azure Database Migration Service (DMS). This functionality limits the amount of downtime that is incurred by the application. References:
https://docs.microsoft.com/en-us/azure/mysql/howto-migrate-online

NEW QUESTION 19
- (Exam Topic 4)
A company has an application that uses Azure SQL Database as the data store.
The application experiences a large increase in activity during the last month of each year.
You need to manually scale the Azure SQL Database instance to account for the increase in data write operations.
Which scaling method should you recommend?

A. Scale up by using elastic pools to distribute resources.
B. Scale out by sharding the data across databases.
C. Scale up by increasing the database throughput units.

Answer: C

Explanation:
As of now, the cost of running an Azure SQL database instance is based on the number of Database Throughput Units (DTUs) allocated for the database. When determining the number of units to allocate for the
solution, a major contributing factor is to identify what processing power is needed to handle the volume of expected requests.
Running the statement to upgrade/downgrade your database takes a matter of seconds.

NEW QUESTION 20
- (Exam Topic 4)
You are designing an Azure Data Factory pipeline for processing data. The pipeline will process data that is stored in general-purpose standard Azure storage.
You need to ensure that the compute environment is created on-demand and removed when the process is completed.
Which type of activity should you recommend?

A. Databricks Python activity
B. Data Lake Analytics U-SQL activity
C. HDInsight Pig activity
D. Databricks Jar activity

Answer: C

Explanation:
The HDInsight Pig activity in a Data Factory pipeline executes Pig queries on your own or on-demand HDInsight cluster.
References:
https://docs.microsoft.com/en-us/azure/data-factory/transform-data-using-hadoop-pig

NEW QUESTION 25
- (Exam Topic 4)
You have an on-premises data warehouse that includes the following fact tables. Both tables have the following columns: DataKey, ProductKey, RegionKey. There are 120 unique product keys and 65 unique region keys.

Table	Comments
Sales	The table is 600 GB in size. Datekey is used extensively in the WHERE clause in queries. ProductKey is used extensively in join operations. RegionKey is used for grouping. Seventy-five percent of records relate to one of 40 regions.
Invoice	The table is 6 GB in size. DataKey and ProductKey are used extensively in the WHERE clause in queries. RegionKey is used for grouping.

Queries that use the data warehouse take a long time to complete.
You plan to migrate the solution to use Azure SQL Data Warehouse. You need to ensure that the Azure-based solution optimizes query performance and minimizes processing skew.
What should you recommend? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

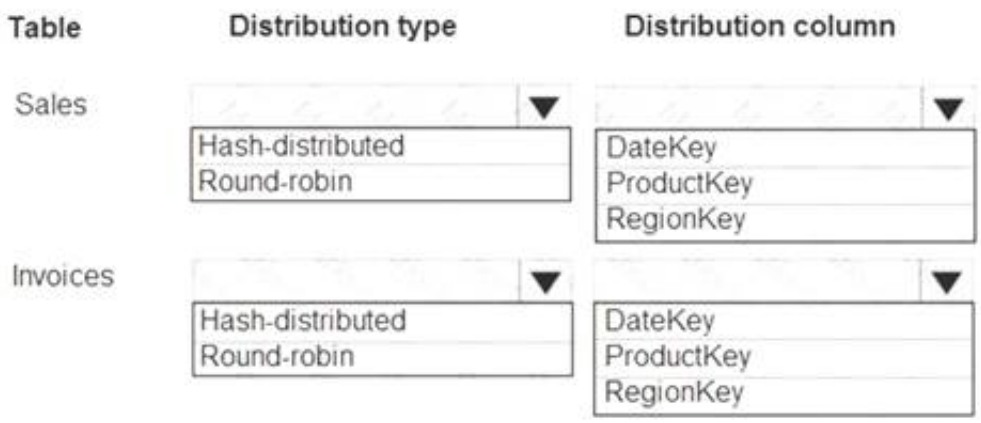

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Hash-distributed
Box 2: ProductKey
ProductKey is used extensively in joins.
Hash-distributed tables improve query performance on large fact tables. Box 3: Round-robin
Box 4: RegionKey
Round-robin tables are useful for improving loading speed.
Consider using the round-robin distribution for your table in the following scenarios:

- When getting started as a simple starting point since it is the default
- If there is no obvious joining key
- If there is not good candidate column for hash distributing the table
- If the table does not share a common join key with other tables
- If the join is less significant than other joins in the query
- When the table is a temporary staging table

Note: A distributed table appears as a single table, but the rows are actually stored across 60 distributions. The rows are distributed with a hash or round-robin algorithm.
References:
https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-distribute

NEW QUESTION 26
- (Exam Topic 4)
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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are designing an Azure SQL Database that will use elastic pools. You plan to store data about customers in a table. Each record uses a value for CustomerID.
You need to recommend a strategy to partition data based on values in CustomerID. Proposed Solution: Separate data into customer regions by using horizontal partitioning. Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
We should use Horizontal Partitioning through Sharding, not divide through regions.
Note: Horizontal Partitioning - Sharding: Data is partitioned horizontally to distribute rows across a scaled out data tier. With this approach, the schema is identical on all participating databases. This approach is also called “sharding”. Sharding can be performed and managed using (1) the elastic database tools libraries or (2)
self-sharding. An elastic query is used to query or compile reports across many shards.
References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-query-overview

NEW QUESTION 28
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your DP-201 Exam with Our Prep Materials Via below:

https://www.certleader.com/DP-201-dumps.html