DP-900 Q&As

Microsoft Azure Data Fundamentals

Pass Microsoft DP-900 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/dp-900.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft Official Exam Center

QUESTION 1

HOTSPOT

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

Statements	Yes	No
A job that processes sales data once daily is an example of a batch workload.	○	○
A job that calculates a rolling average temperature reading is an example of a streaming workload.	○	○
A job that calculates average revenue per product for the last month is an example of a streaming workload.	○	○

Correct Answer:

Statements	Yes	No
A job that processes sales data once daily is an example of a batch workload.	●	○
A job that calculates a rolling average temperature reading is an example of a streaming workload.	●	○
A job that calculates average revenue per product for the last month is an example of a streaming workload.	○	●

Box 1: Yes Batch processing refers to processing of high volume of data in batch within a specific time span. It processes large volume of data all at once. Batch processing is used when data size is known and finite. It takes little longer time to processes data.

Box 2: Yes

Stream processing refers to processing of continuous stream of data immediately as it is produced. It analyzes streaming data in real time. Stream processing is used when the data size is unknown and infinite and continuous.

Box 3: No

Reference:

https://www.geeksforgeeks.org/difference-between-batch-processing-and-stream-processing/

QUESTION 2

To complete the sentence, select the appropriate option in the answer area.

Hot Area:

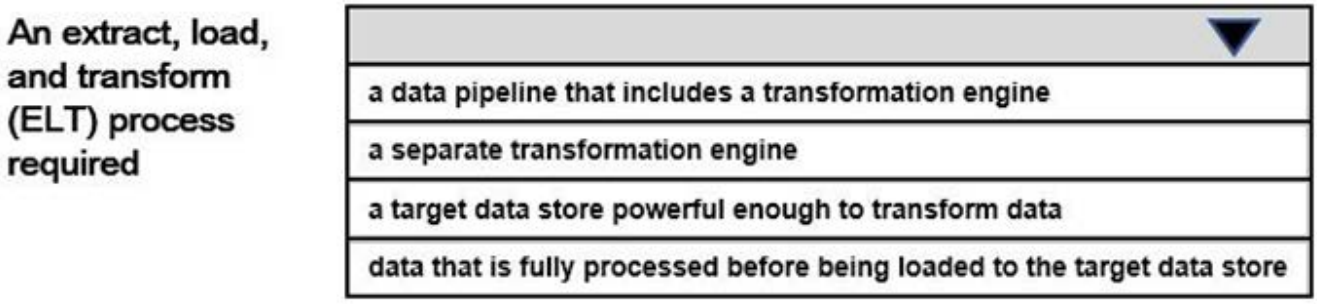

Correct Answer:

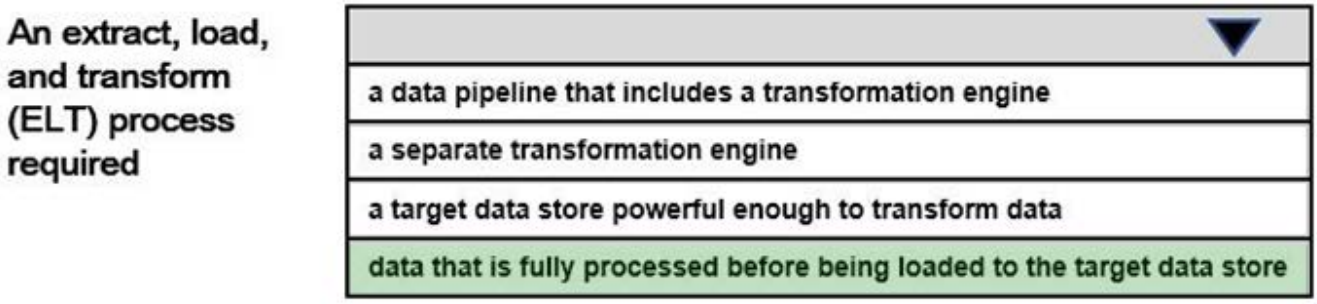

QUESTION 3

HOTSPOT

For each of the following statements, select Yes if the statement is true, Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

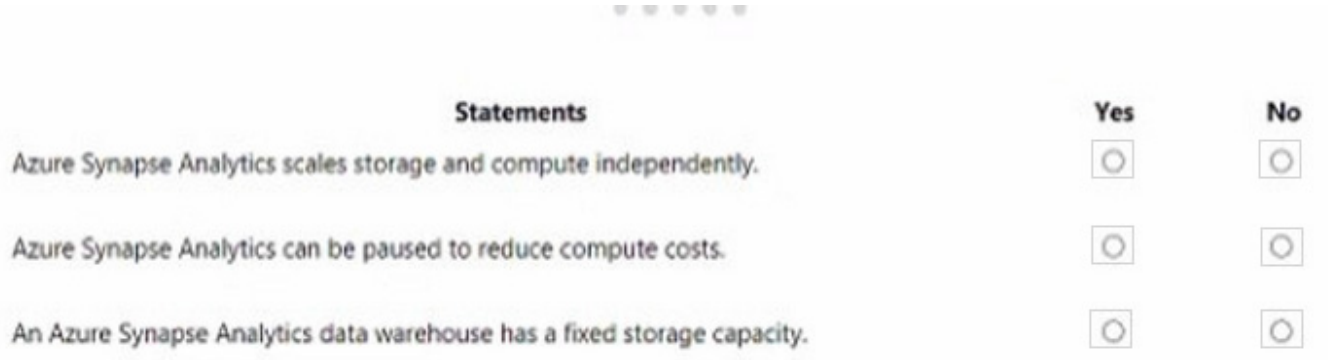

Correct Answer:

QUESTION 4

A bar chart showing year-to-date sales by region is an example of which type of analytics?

A. Predictive

B. Prescriptive

C. Diagnostic

D. Descriptive

Correct Answer: D

QUESTION 5

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

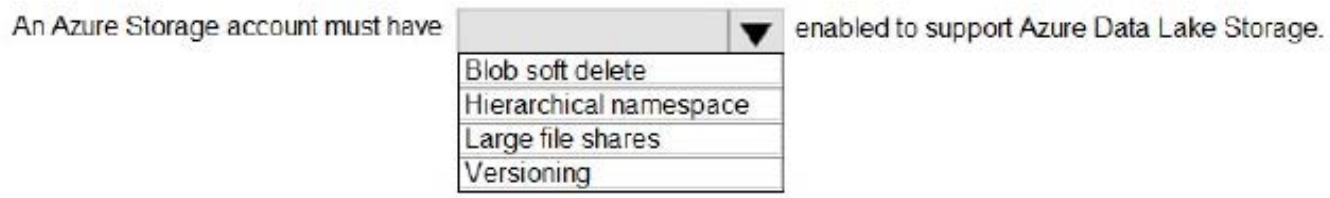

Correct Answer:

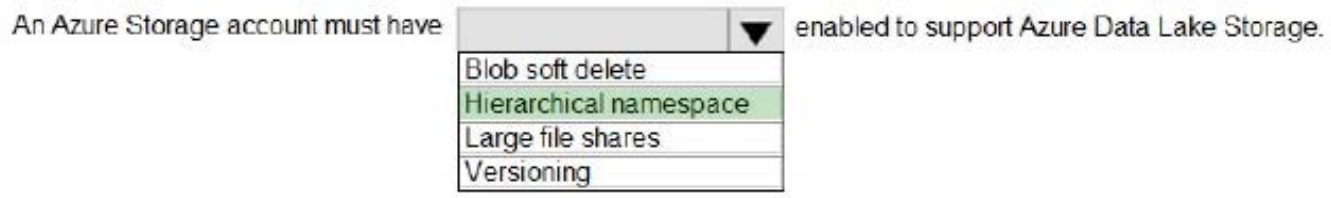

QUESTION 6

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

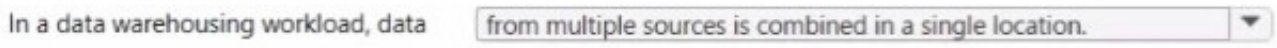

Correct Answer:

QUESTION 7

HOTSPOT

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Hot Area:

Statements	Yes	No
Data stored in the Hot access tier of Azure Blob storage is stored on high-performance media.	○	○
The Cool access tier in Azure Blob storage incurs higher storage costs compared to the Hot access tier.	○	○
The Archive access tier in Azure Blob storage has the highest latency of the three available tiers.	○	○

Correct Answer:

Statements	Yes	No
Data stored in the Hot access tier of Azure Blob storage is stored on high-performance media.	○	○
The Cool access tier in Azure Blob storage incurs higher storage costs compared to the Hot access tier.	○	○
The Archive access tier in Azure Blob storage has the highest latency of the three available tiers.	○	○

QUESTION 8

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

In a relational database, each row in a table has

- a different set of columns
- a key-value pair
- the same set of columns
- unstructured data

Correct Answer:

In a relational database, each row in a table has

- a different set of columns
- a key-value pair
- the same set of columns
- unstructured data

QUESTION 9

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

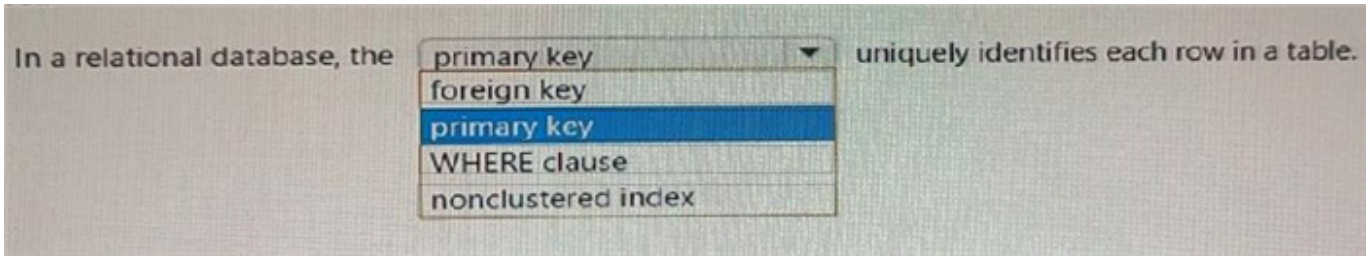

Correct Answer:

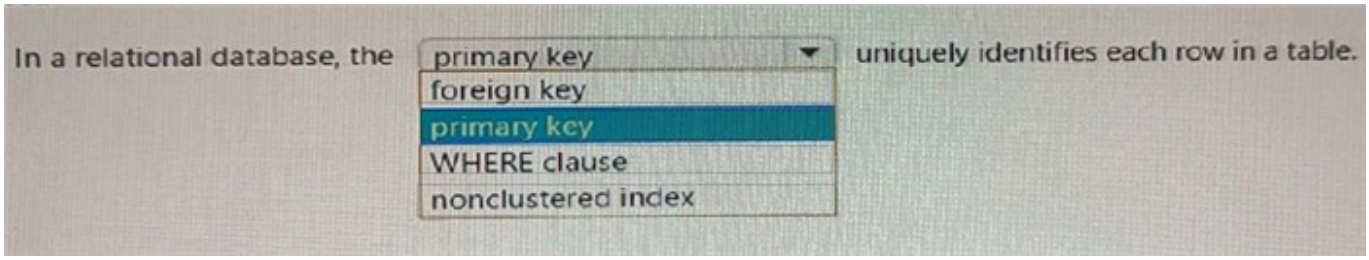

QUESTION 10

Which storage solution supports access control lists (ACLs) at the file and folder level?

A. Azure Data Lake Storage

B. Azure Queue storage

C. Azure Blob storage

D. Azure Cosmos DB

Correct Answer: A

Azure Data Lake Storage Gen2 implements an access control model that supports both Azure role-based access control (Azure RBAC) and POSIX-like access control lists (ACLs).

Reference: https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-access-control

QUESTION 11

HOTSPOT

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

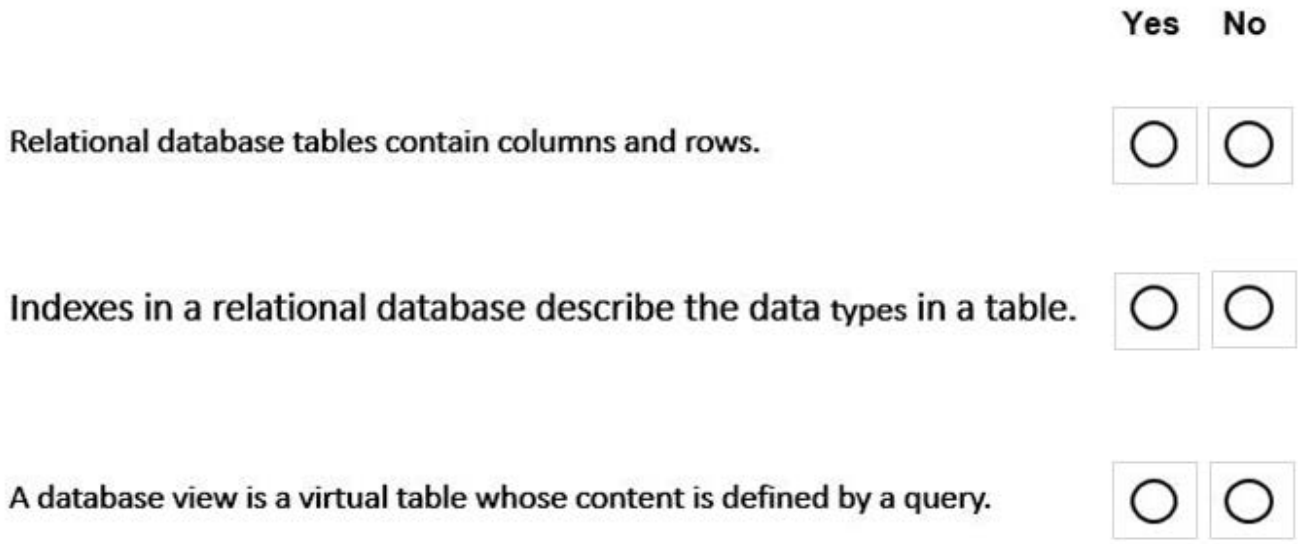

Correct Answer:

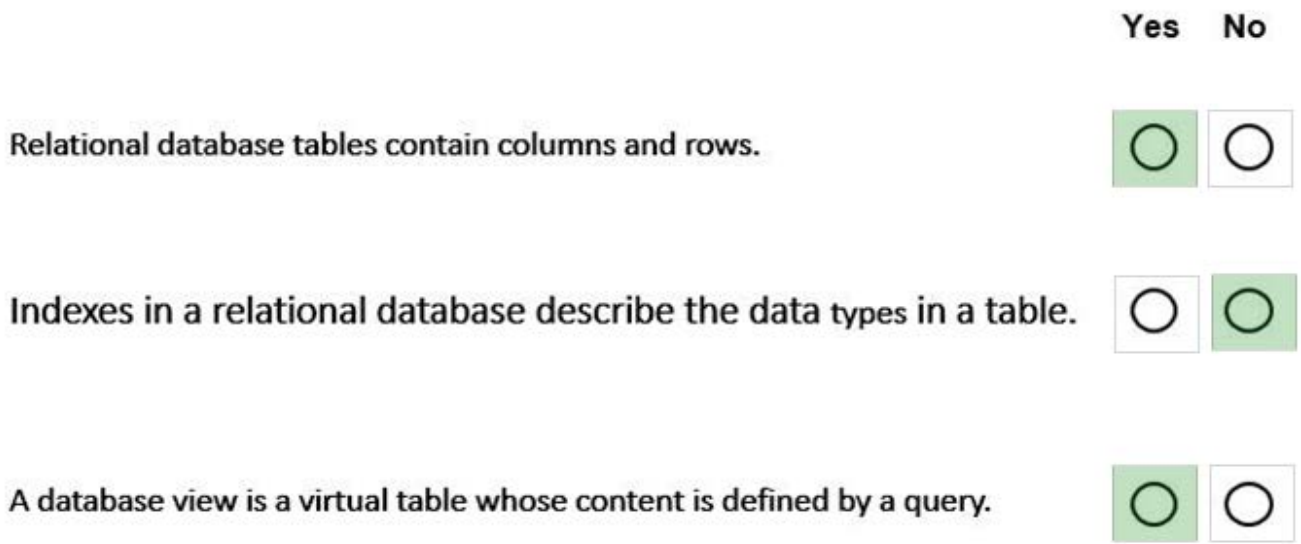

Box 1: Yes

Tables are database objects that contain all the data in a database. In tables, data is logically organized in a row-and-column format similar to a spreadsheet. Each row represents a unique record, and each column represents a field in the

record.

Box 2: No

An index is an on-disk structure associated with a table or view that speeds retrieval of rows from the table or view.

Box 3: Yes

A view is a virtual table whose contents are defined by a query. Like a table, a view consists of a set of named columns and rows of data.

Reference:

https://docs.microsoft.com/en-us/sql/relational-databases/tables/tables

https://docs.microsoft.com/en-us/sql/relational-databases/indexes/clustered-and-nonclustered-indexes-described

https://docs.microsoft.com/en-us/sql/relational-databases/views/views?view=sql-server-ver15

QUESTION 12

Your company needs to design a database that illustrates the relationships between utilization levels of individual network devices across a local area network. Which type of data store should you use?

A. Key/value

B. Graph

C. Document

D. columnar

Correct Answer: B

Data as it appears in the real world is naturally connected. Traditional data modeling focuses on defining entities separately and computing their relationships at runtime. While this model has its advantages, highly connected data can be challenging to manage under its constraints.

A graph database approach relies on persisting relationships in the storage layer instead, which leads to highly efficient graph retrieval operations. Azure Cosmos DB\\'s Gremlin API supports the property graph model.

Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/graph-introduction#introduction-to-graph-databases

---

QUESTION 13

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

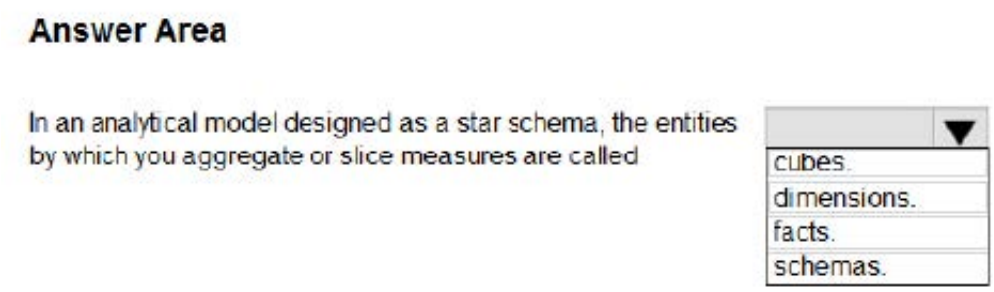

Correct Answer:

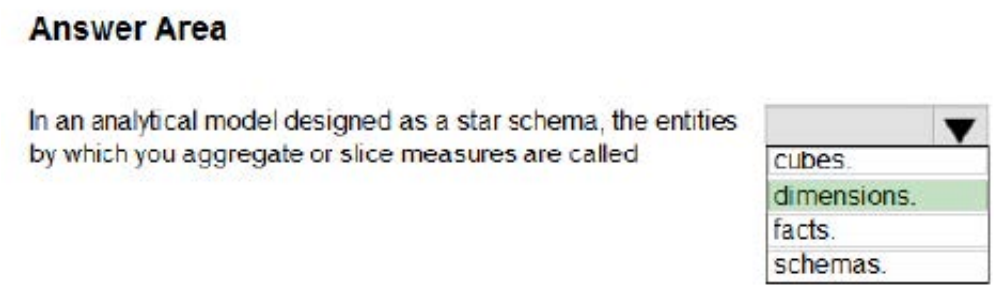

Dimensions:- In a star schema, the entities by which you aggregate or slice measures are called dimensions. Dimensions are the descriptive attributes or perspectives of the data that provide the context for the measures (also known as facts)

in the model.

For example, in a sales data star schema, the dimensions could be time, product, store location, and customer. These dimensions provide context for the sales measures such as total sales, units sold, and revenue. By aggregating or slicing

the measures by different dimensions, you can gain insights into the data from different perspectives.

QUESTION 14

HOTSPOT

To complete the sentence, select the appropriate option in the answer area.

Hot Area:

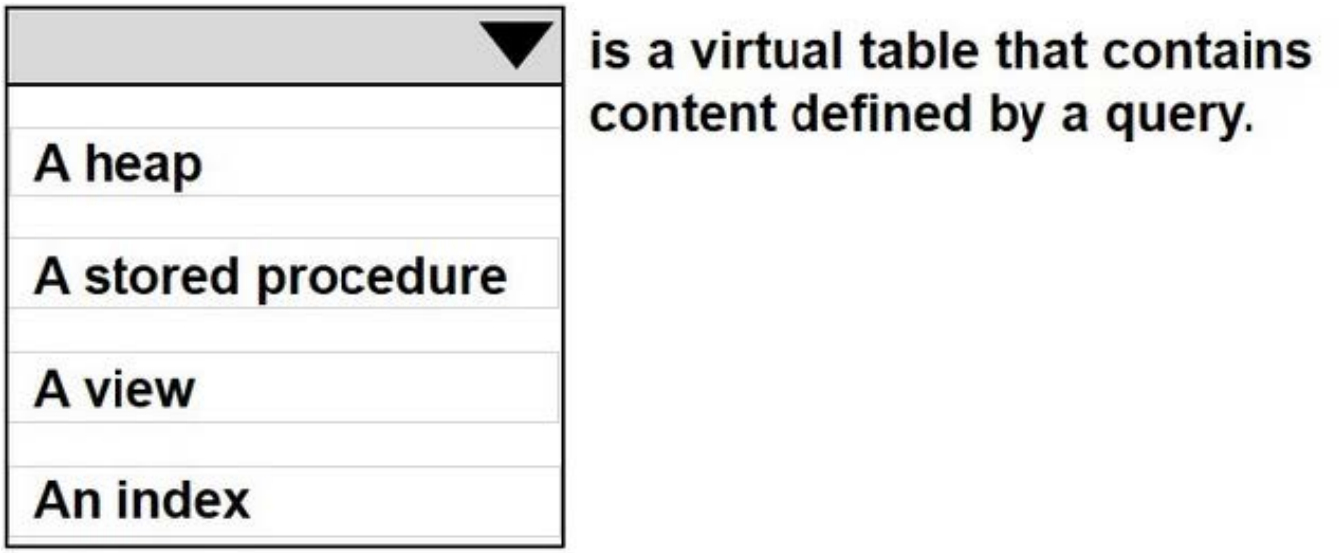

Correct Answer:

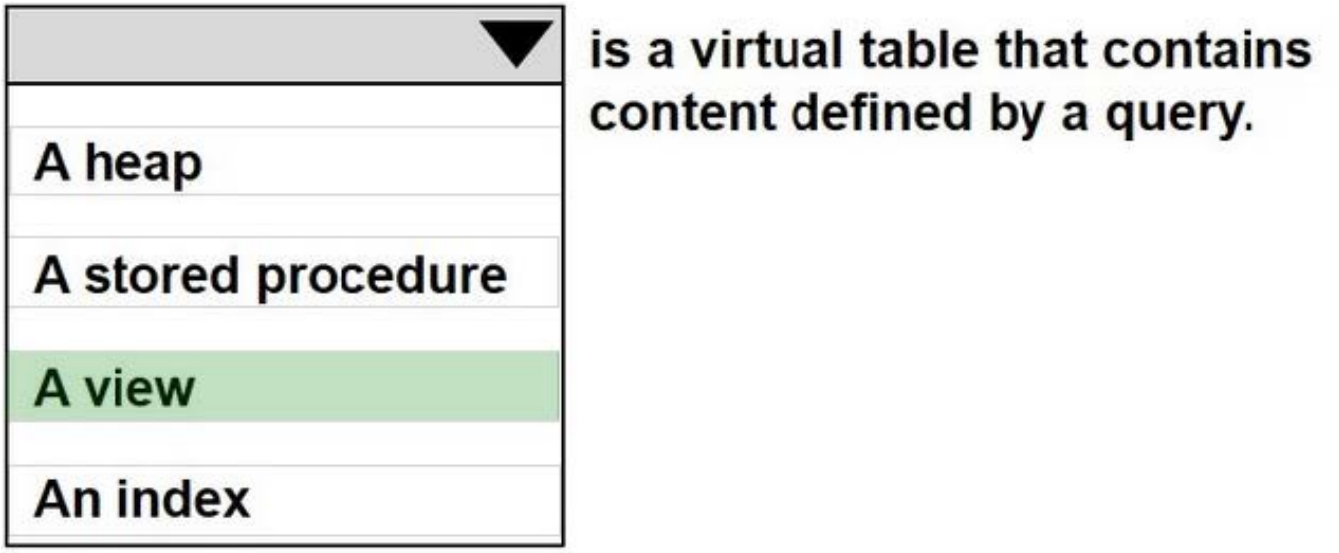

Reference: https://docs.microsoft.com/en-us/sql/relational-databases/views/views

QUESTION 15

HOTSPOT

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

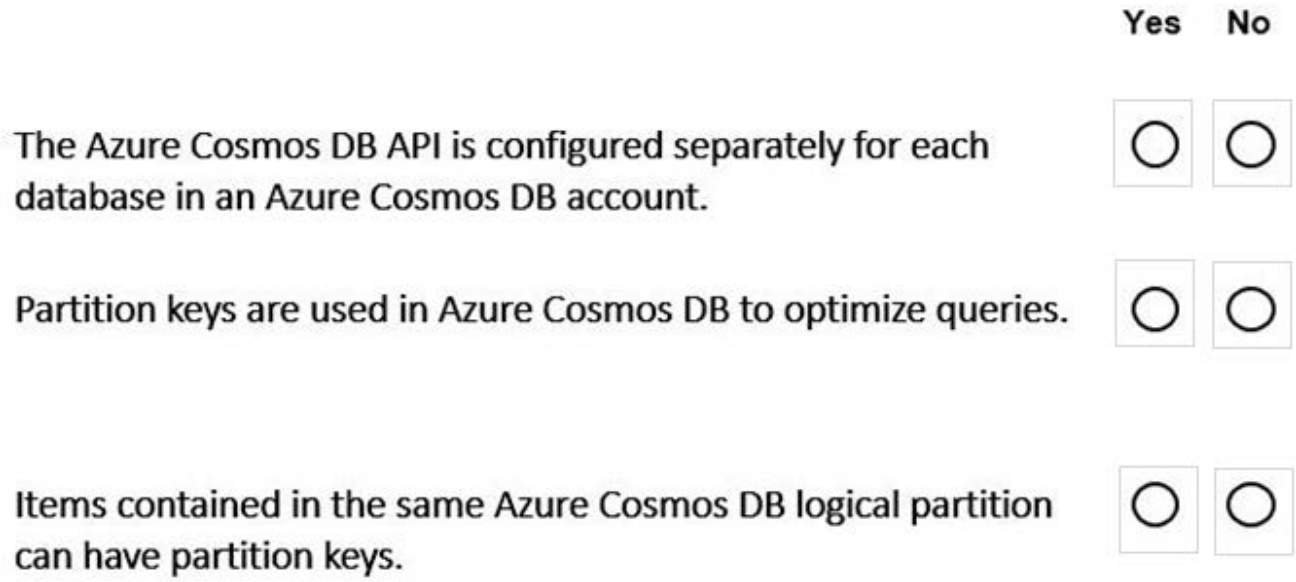

Correct Answer:

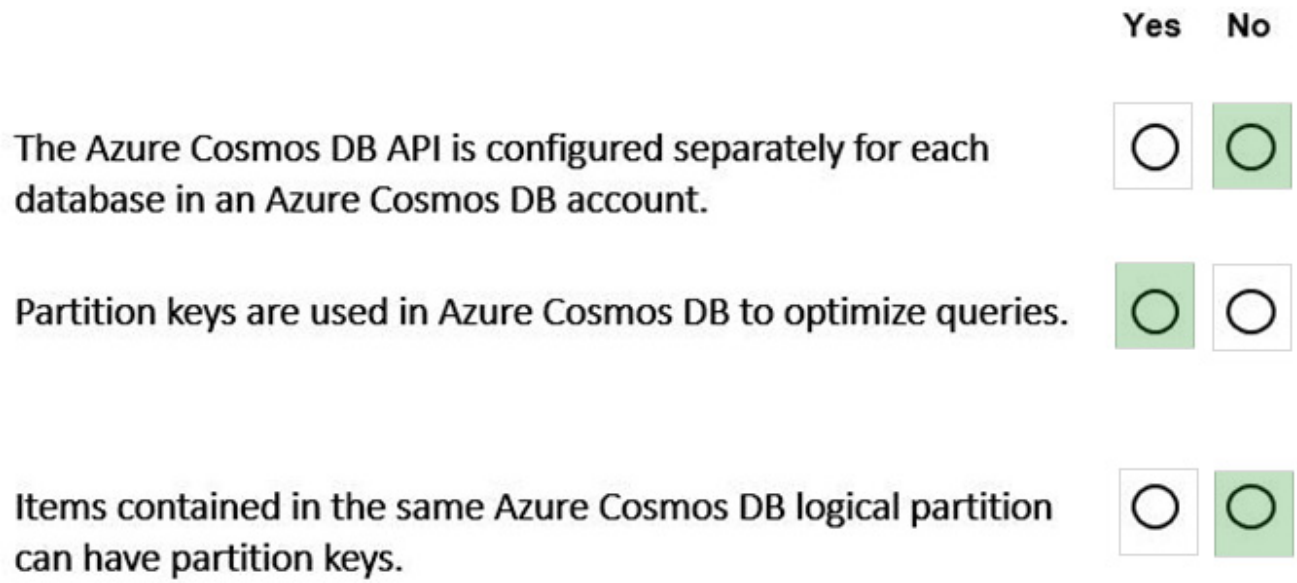

Box 1: No

The API determines the type of account to create. Azure Cosmos DB provides five APIs: Core (SQL) and MongoDB for document data, Gremlin for graph data, Azure Table, and Cassandra. Currently, you must create a separate account for

each API.

Box 2: Yes

Azure Cosmos DB uses partitioning to scale individual containers in a database to meet the performance needs of your application. In partitioning, the items in a container are divided into distinct subsets called logical partitions. Logical

partitions are formed based on the value of a partition key that is associated with each item in a container.

Box 3: No

Logical partitions are formed based on the value of a partition key that is associated with each item in a container.

Reference:

https://docs.microsoft.com/en-us/azure/cosmos-db/partitioning-overview