Microsoft

(DP-900)

Microsoft Azure Data Fundamentals

Total: **302 Questions**

Link:

Question: 1

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

Answer Area

Descriptive analytics tells you

- what is most likely to occur in the future.
- what occurred in the past.
- which actions you can perform to affect outcomes.
- why something occurred in the past.

Answer:

Answer Area

Descriptive analytics tells you

- what is most likely to occur in the future.
- what occurred in the past.
- which actions you can perform to affect outcomes.
- why something occurred in the past.

Explanation:

Descriptive analytics tells "What happened in the past". So B is the right choice.

Reference:

https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitive-analytics/

Question: 2

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:

Answer Area

Statements	Yes	No
Normalization involves eliminating relationships between database tables.	○	○
Normalizing a database reduces data redundancy.	○	○
Normalization improves data integrity.	○	○

Answer:

Answer Area

Statements	Yes	No
Normalization involves eliminating relationships between database tables.	○	○
Normalizing a database reduces data redundancy.	○	○
Normalization improves data integrity.	○	○

Explanation:
Normalization is the process of organizing data in a database. This includes creating tables and establishing relationships between those tables according to rules designed both to protect the data and to make the database more flexible by eliminating redundancy and inconsistent dependency.

Reference:
https://docs.microsoft.com/en-us/office/troubleshoot/access/database-normalization-description#:~:text=Normalization%20is%20the%20process%20of,eliminating%20redundancy%20and%20inconsistent%20dependency
https://www.sqlshack.com/what-is-database-normalization-in-sql-server/

Question: 3 HOTSPOT
-
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

Answer Area

An extract, transform, and load (ETL) process requires

- a matching schema in the data source and the data target.
- a target data store powerful enough to transform data.
- data that is fully processed before being loaded to the target data store.
- that the data target be a relational database.

Answer:

Answer Area

An extract, transform, and load (ETL) process requires

- a matching schema in the data source and the data target.
- a target data store powerful enough to transform data.
- data that is fully processed before being loaded to the target data store.
- that the data target be a relational database.

Explanation:
Extract, transform, and load (ETL) is a data pipeline used to collect data from various sources, transform the data according to business rules, and load it into a destination data store.

Reference:
https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl

Question: 4

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

Answer Area

In batch processing,

- data is always inserted one row at a time.
- data is processed in real-time.
- latency in delivering data processing results is acceptable.
- processing can only execute serially.

Answer:

Answer Area

In batch processing, [▼]

- data is always inserted one row at a time.
- data is processed in real-time.
- latency in delivering data processing results is acceptable.
- processing can only execute serially.

Explanation:
When to use batch processing.

You might expect latencies when using batch processing. For many situations, however, this type of delay before the transfer of data begins is not a big issue - the processes that use this function are not mission critical at the exact moment.

Reference:
https://www.bmc.com/blogs/what-is-batch-processing-batch-processing-explained/

Question: 5

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

Answer Area

Transcribing audio files is an example of [▼] analytics.

- cognitive
- descriptive
- predictive
- prescriptive

Answer:

Answer Area

Transcribing audio files is an example of [▼] analytics.

- cognitive
- descriptive
- predictive
- prescriptive

Explanation:

Transcribe means converting from Audio to Text. This is similar to Alexa, Alexa is a Cognitive device. Cognitive is the correct answer.

Reference:
https://azure.microsoft.com/en-us/services/cognitive-services/speech-services/

Question: 6 DRAG

DROP -
Match the types of analytics that can be used to answer the business questions.
To answer, drag the appropriate analytics type from the column on the left to its question on the right. Each analytics type may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.
Select and Place:

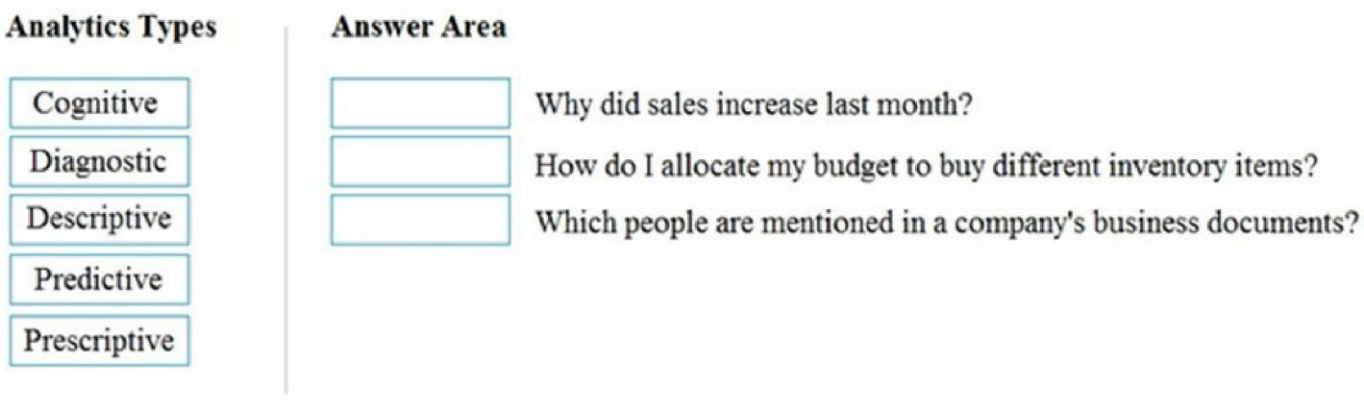

Answer:

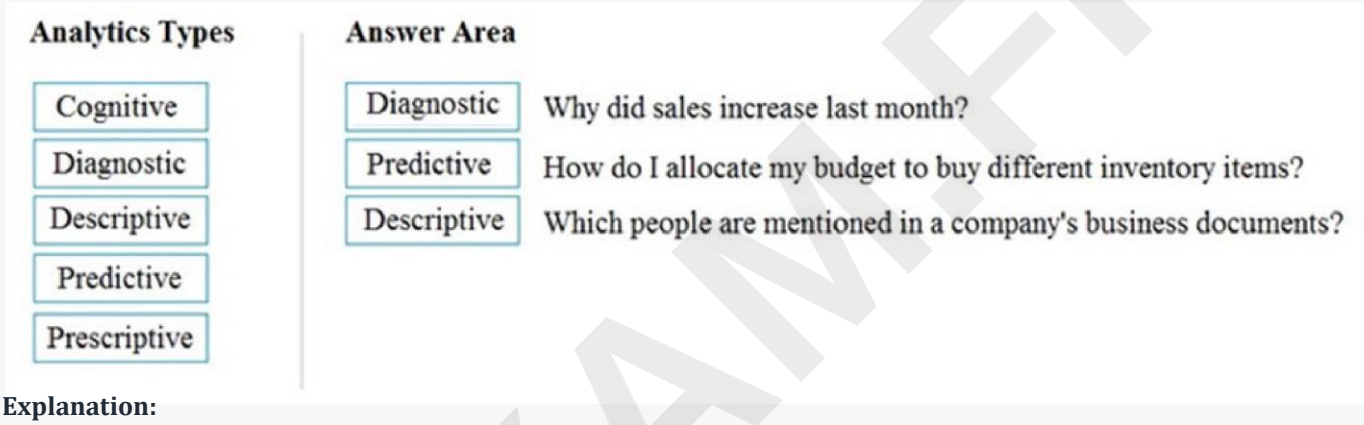

Explanation:
Box 1: Diagnostic -
Diagnostic Analytics: At this stage you can begin to answer some of those why questions. Historical data can begin to be measured against other data to answer the question of why something happened in the past. This is the process of gathering and interpreting different data sets to identify anomalies, detect patters, and determine relationships.

Box 2: Prescriptive -
Prescriptive analytics is a combination of data, mathematical models, and various business rules to infer actions to influence future desired outcomes.

Predictive analytics, broadly speaking, is a category of business intelligence that uses descriptive and

predictive variables from the past to analyze and identify the likelihood of an unknown future outcome Box 3: Descriptive -

☞ Generally speaking, data analytics comes in four types:

☞ Descriptive, to answer the question: What's happening?

☞ Diagnostic, to answer the question: Why's happening?

☞ Predictive, to answer the question: What will happen?

☞ Prescriptive, to answer the question: What actions should we take?

Reference: https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitive-analytics/ https://azure.microsoft.com/en-us/blog/answering-whats-happening-whys-happening-and-what-will-happen-with-iot-analytics/

Question: 7 HOTSPOT -

You have the following JSON document.

```
"customer" : {
    "first name" : "Ben",
    "last name" : "Smith",
    "address" : {
        "line 1" : "161 Azure Ln",
        "line 2" : "Palo Alto",
        "ZIP code" : "54762"
    },
    "social media": [
        {
            "service" : "twitter",
            "handle" : "@bensmith"
        },
        {
            "service" : "linkedin",
            "handle" : "bensmith"
        }
    ],
    "phone numbers": [
        {
            "type" : "mobile",
            "number" : "555-555-555"
        }
    ]
}
```

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the JSON document.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

Customer is **[answer choice]**.

- a nested array
- a nested object
- a root object

Address is **[answer choice]**.

- a nested array
- a nested object
- a root object

Social media is **[answer choice]**.

- a nested array
- a nested object
- a root object

Answer:

Answer Area

Customer is **[answer choice]**.

- a nested array
- a nested object
- **a root object**

Address is **[answer choice]**.

- a nested array
- **a nested object**
- a root object

Social media is **[answer choice]**.

- **a nested array**
- a nested object
- a root object

Explanation:

Reference:
https://www.w3schools.com/js/js_json_arrays.asp
https://www.w3schools.com/js/js_json_objects.asp

Question: 8

HOTSPOT -
You are reviewing the data model shown in the following exhibit.

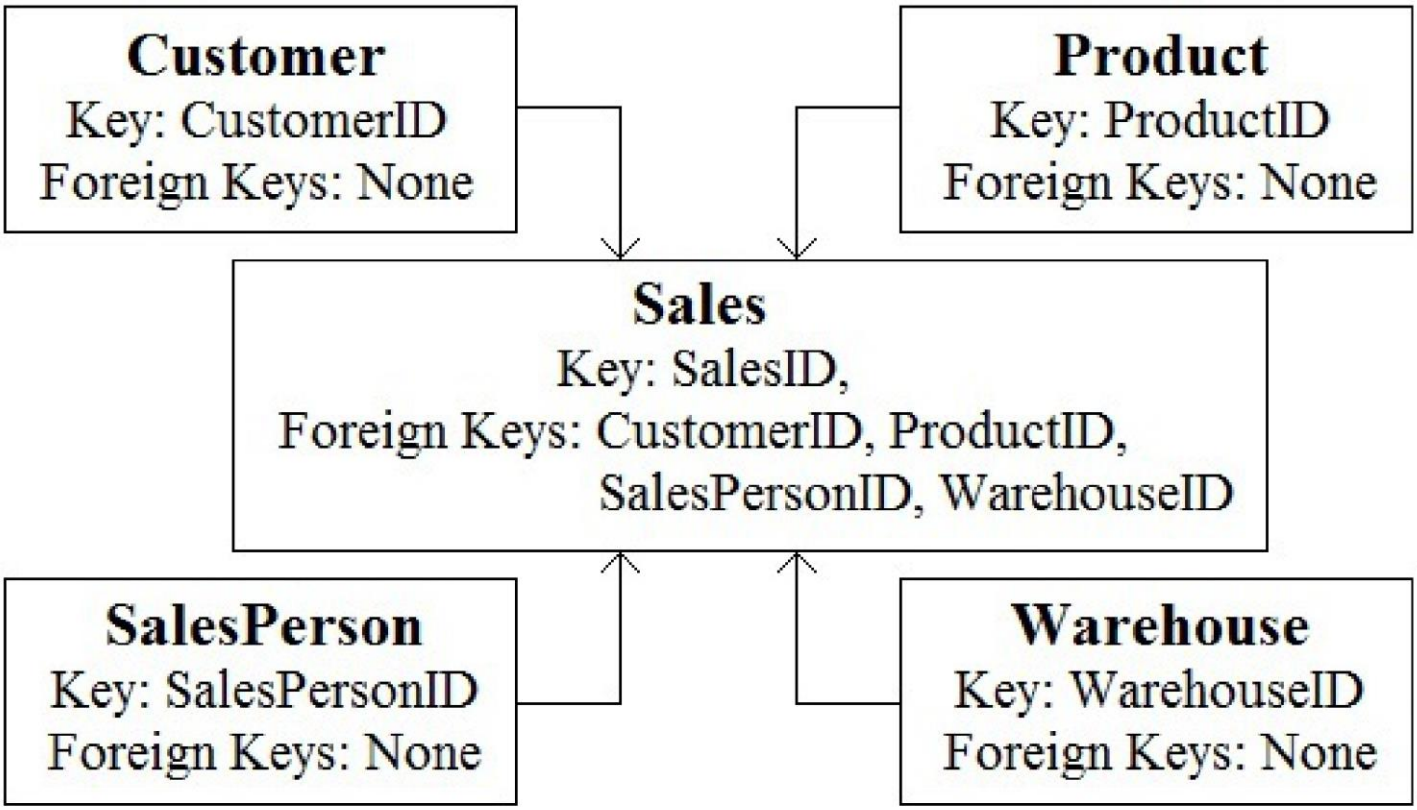

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point
Hot Area:

Answer Area

The data model is a **[answer choice]**.

- transactional model
- star schema
- snowflake schema

Customer is a **[answer choice]** table.

- fact
- dimension
- bridge

Answer:

Answer Area

The data model is a **[answer choice]**.

- transactional model
- star schema
- snowflake schema

Customer is a **[answer choice]** table.

- fact
- dimension
- bridge

Explanation:

Box 1: star schema -

In computing, the star schema is the simplest style of data mart schema and is the approach most widely used to develop data warehouses and dimensional data marts. The star schema consists of one or more fact tables referencing any number of dimension tables. The star schema is an important special case of the snowflake schema, and is more effective for handling simpler queries.

Example:

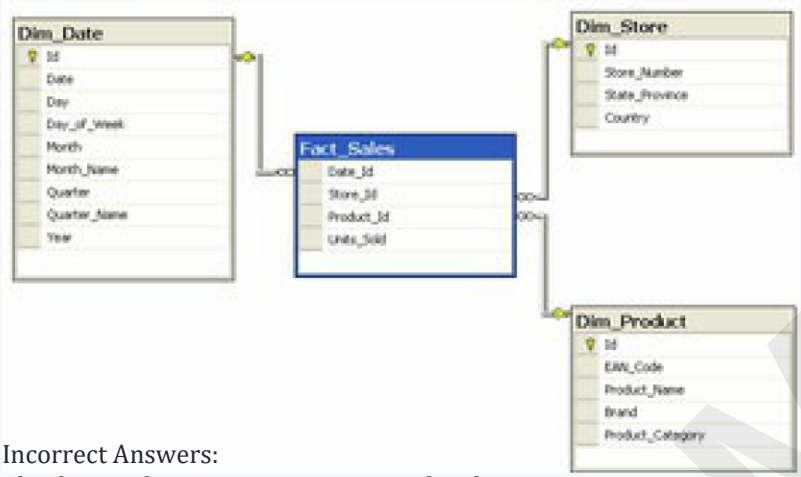

Incorrect Answers:

The data in the question is not normalized.

The snowflake schema is a variation of the star schema, featuring normalization of dimension tables. Example:

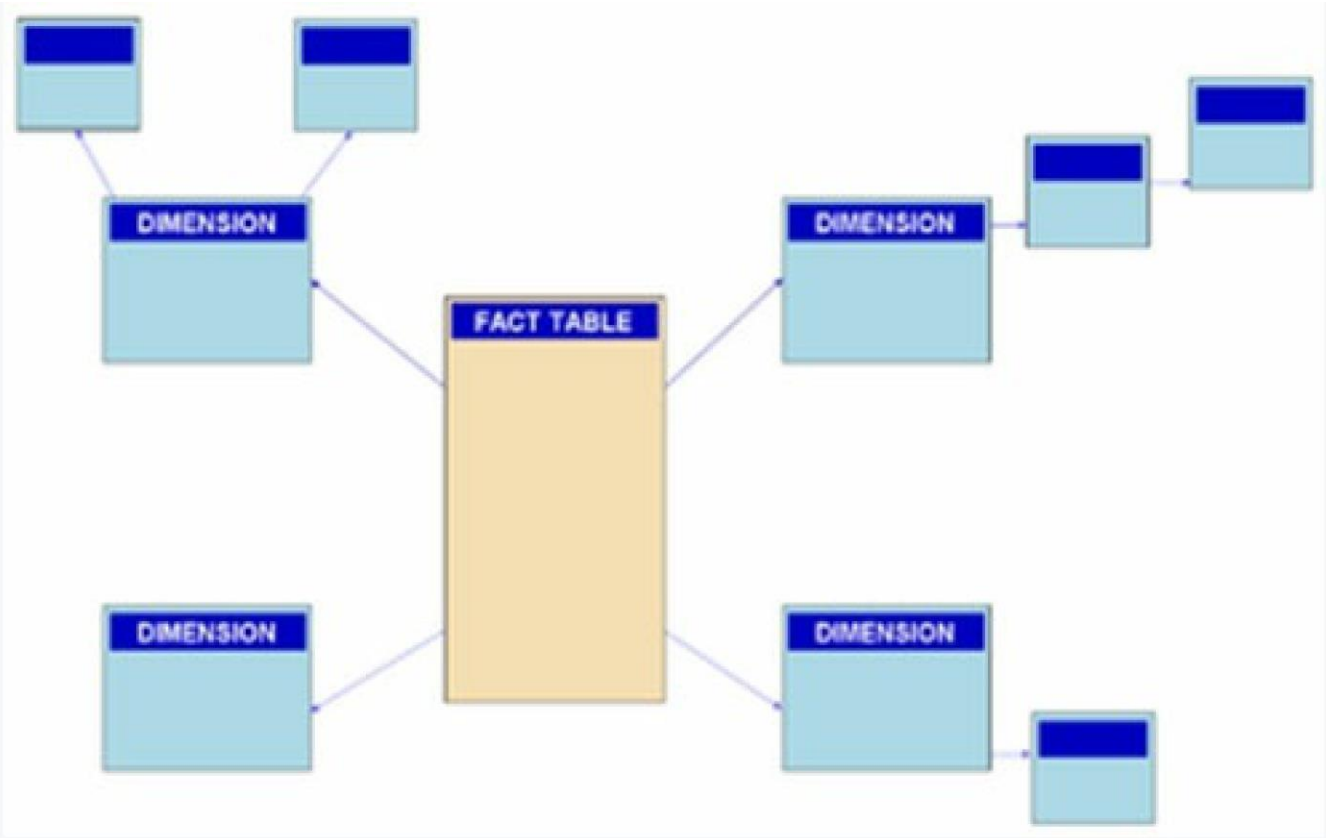

Note: A snowflake schema is a logical arrangement of tables in a multidimensional database such that the entity relationship diagram resembles a snowflake shape. The snowflake schema is represented by centralized fact tables which are connected to multiple dimensions.[citation needed]. "Snowflaking" is a method of normalizing the dimension tables in a star schema. When it is completely normalized along all the dimension tables, the resultant structure resembles a snowflake with the fact table in the middle.

Box 2: dimension -
The star schema consists of one or more fact tables referencing any number of dimension tables.

Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-o verview https://en.wikipedia.org/wiki/Star_schema https://en.wikipedia.org/wiki/Snowflake_schema https://az ure.microsoft.com/en-us/blog/data-models-within-azure-analysis-services-and-power-bi/

Question: 9

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

Answer Area

The massively parallel processing (MPP) engine of Azure Synapse Analytics

- distributes processing across compute nodes.
- distributes processing across control nodes.
- redirects client connections across compute nodes.
- redirects client connections across control nodes.

Answer:

Answer Area

The massively parallel processing (MPP) engine of Azure Synapse Analytics

- distributes processing across compute nodes.
- distributes processing across control nodes.
- redirects client connections across compute nodes.
- redirects client connections across control nodes.

Explanation:

Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/massively-parallel-processin g-mpp-architecture

Question: 10

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

Answer Area

- A clustered index
- A FileTable
- A foreign key
- A stored procedure

is an object associated with a table that sorts and stores the data rows in the table based on their key values.

Answer:

Answer Area

- A clustered index
- A FileTable
- A foreign key
- A stored procedure

is an object associated with a table that sorts and stores the data rows in the table based on their key values.

Explanation:

Clustered indexes sort and store the data rows in the table or view based on their key values. These are the columns included in the index definition. There can be only one clustered index per table, because the data rows themselves can be stored in only one order.

Reference:

https://docs.microsoft.com/en-us/sql/relational-databases/indexes/clustered-and-nonclustered-indexes-described?view=sql-server-ver15

Question: 11

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

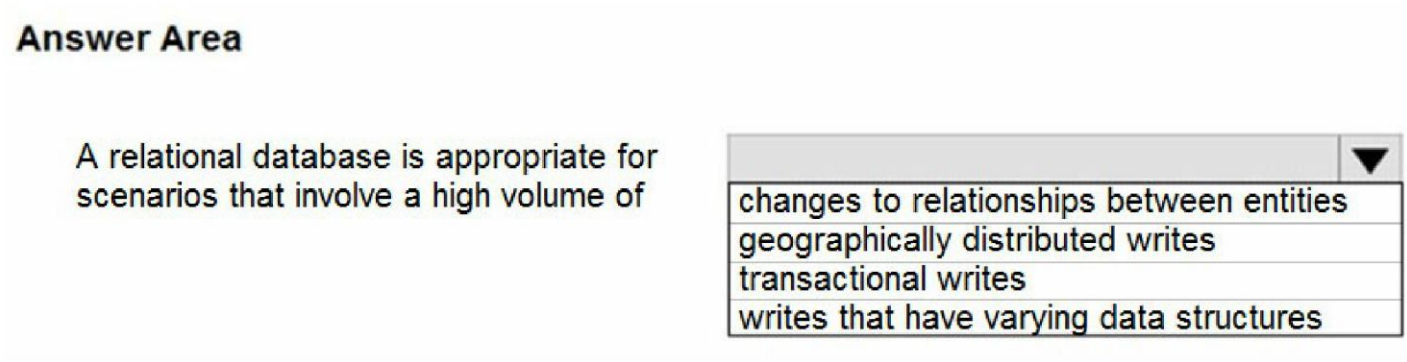

Answer:

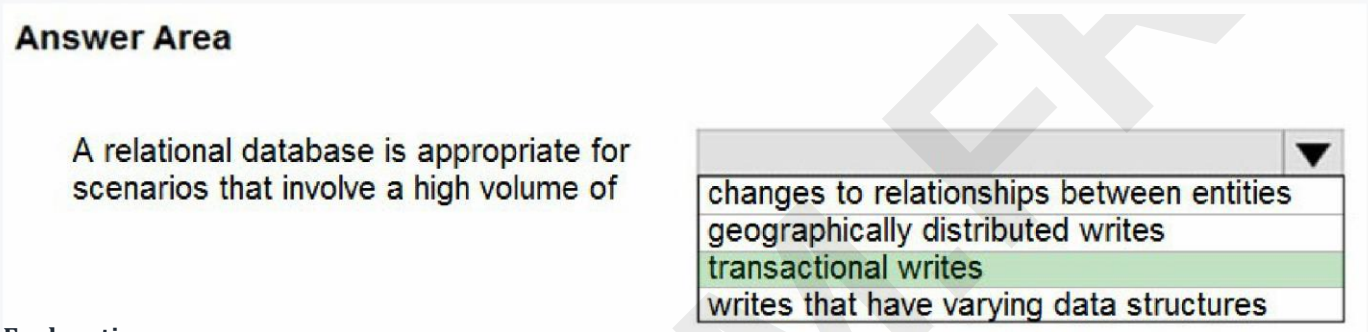

Explanation:
Disadvantages of non-relational databases include: Data Consistency " non-relational databases do not perform ACID transactions.

Note: Relational databases are optimized for writes. They are optimized for consistency and availability. Advantages of relational databases include simplicity, ease of data retrieval, data integrity, and flexibility.

Incorrect Answers:
Use a relational database when data that you work with is structured, and the structure is not subject to frequent changes.

Use Cloud storage (no relational database) for geographically distributed writes.

Reference:
https://towardsdatascience.com/choosing-the-right-database-c45cd3a28f77

Question: 12

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:

Answer Area

Statements	Yes	No
Batch processing can output data to a file store	○	○
Batch processing can output data to a relational database	○	○
Batch processing can output data to a NoSQL database	○	○

Answer:

Answer Area

Statements	Yes	No
Batch processing can output data to a file store	○	○
Batch processing can output data to a relational database	○	○
Batch processing can output data to a NoSQL database	○	○

Explanation:

Batch processing can output data to a file store:

Correct Answer: Yes

Explanation: Batch processing can indeed output data to a file store. This is commonly used for storing large volumes of data in files, such as logs, reports, or backups.

Batch processing can output data to a relational database:

Correct Answer: Yes

Explanation: Batch processing can also output data to a relational database. This is useful for scenarios where structured data needs to be stored and queried efficiently.

Batch processing can output data to a NoSQL database:

Correct Answer: Yes

Explanation: Yes, batch processing can also output data to a NoSQL database. NoSQL databases are designed to handle unstructured or semi-structured data, making them suitable for various use cases.

Reference:

https://docs.microsoft.com/en-us/azure/architecture/data-guide/big-data/batch-processing

Question: 13

DRAG DROP -
Your company plans to load data from a customer relationship management (CRM) system to a data warehouse by using an extract, load, and transform (ELT) process.
Where does data processing occur for each stage of the ELT process? To answer, drag the appropriate locations to the correct stages. Each location may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.
Select and Place:

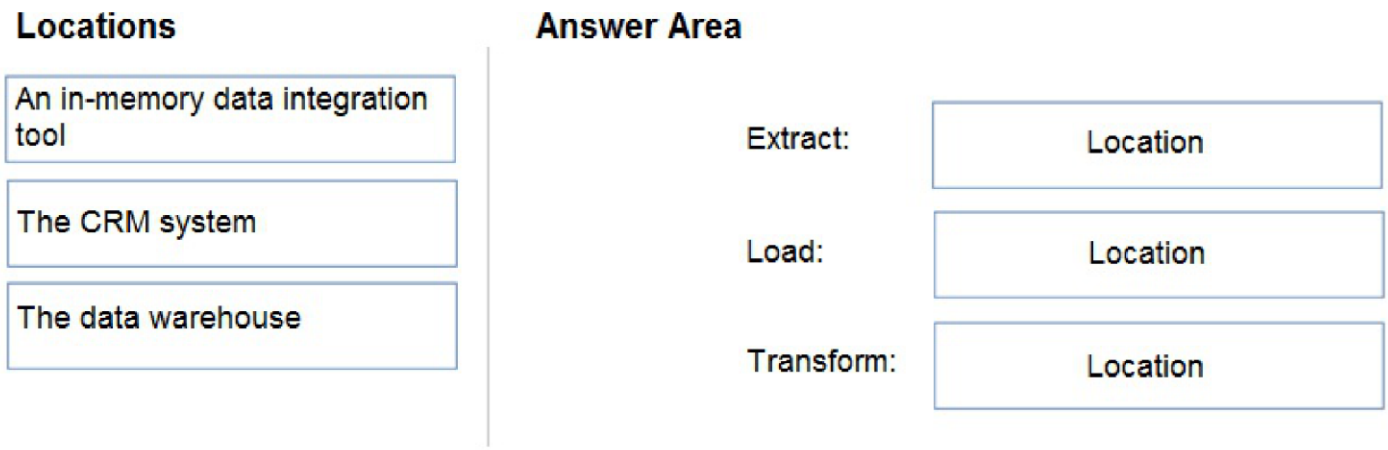

Answer:

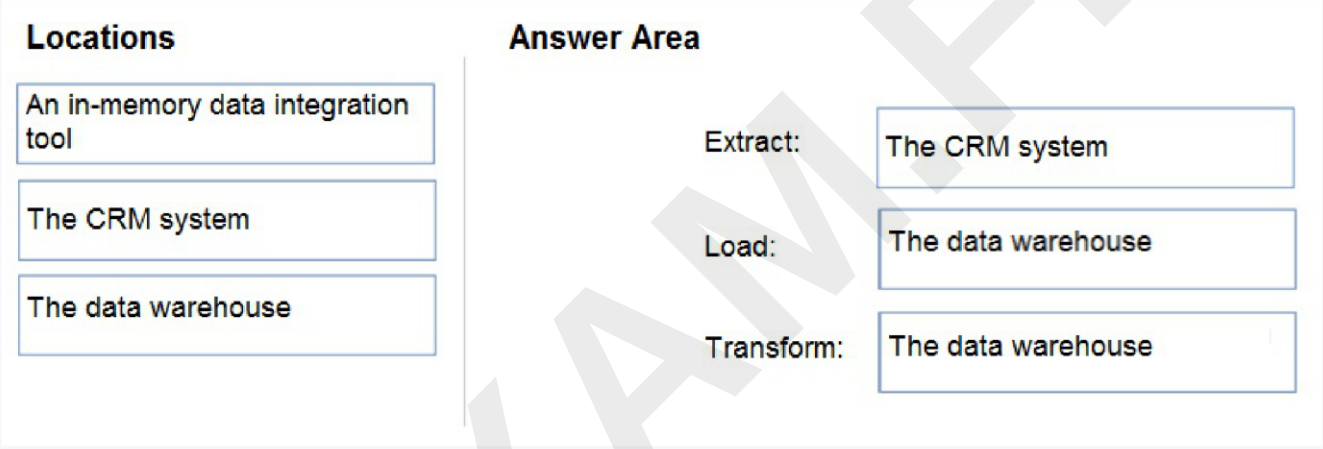

Explanation:

Extract: The CRM System

Explanation: In the ELT process, data is extracted directly from the source system, which in this case is the CRM system.

Load: The Data Warehouse

Explanation: In an ELT process, data is loaded directly into the data warehouse without being transformed first.

Transform: The Data Warehouse

Explanation: In ELT, transformation occurs after data is loaded into the data warehouse, using the data warehouse's processing capabilities to perform the transformations.

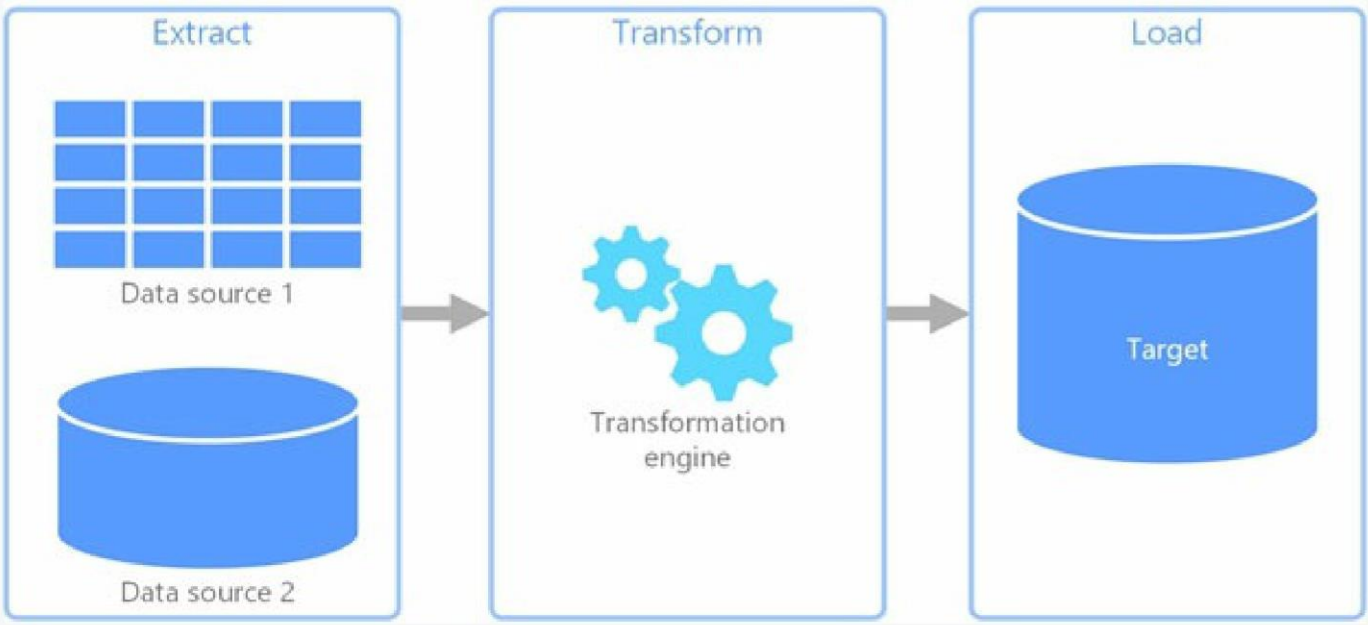

Reference:

https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl

Question: 14 HOTSPOT

-

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

Hot Area:

Answer Area

A visualization that shows a university's current student enrollment versus the maximum capacity is an example of [▼] analytics.

- cognitive
- descriptive
- predictive
- prescriptive

Answer:

Answer Area

A visualization that shows a university's current student enrollment versus the maximum capacity is an example of [▼] analytics.

- cognitive
- descriptive
- predictive
- prescriptive

Explanation:
Generally speaking, data analytics comes in four types (Figure 1):
Descriptive, to answer the question: What's happening?
Diagnostic, to answer the question: Why's happening?
Predictive, to answer the question: What will happen?
Prescriptive, to answer the question: What actions should we take?

Reference:
https://azure.microsoft.com/en-us/blog/answering-whats-happening-whys-happening-and-what-will-happen-with-iot-analytics/

Question: 15

DRAG DROP -
Match the types of visualizations to the appropriate descriptions.
To answer, drag the appropriate visualization type from the column on the left to its description on the right. Each visualization type may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.
Select and Place:

Visualizations	Answer Area	
Key influencer	Visualization	A chart of colored, nested rectangles that displays individual data points represented by the size and color of a relative rectangle.
Scatter	Visualization	A chart that displays the major contributors of a selected result or value.
Treemap	Visualization	A chart that shows the relationship between two numerical values.

Answer:

Visualizations	Answer Area	
Key influencer	Treemap	A chart of colored, nested rectangles that displays individual data points represented by the size and color of a relative rectangle.
Scatter	Key influencer	A chart that displays the major contributors of a selected result or value.
Treemap	Scatter	A chart that shows the relationship between two numerical values.

Explanation:
Box 1: Tree map -
Treemaps are charts of colored rectangles, with size representing value. They can be hierarchical, with

rectangles nested within the main rectangles.

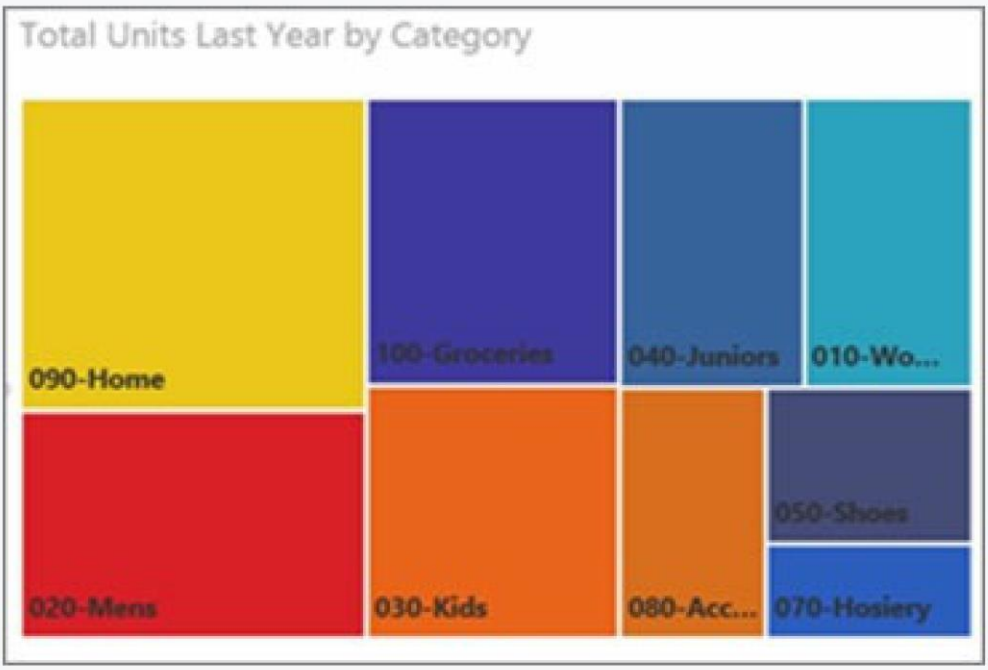

Box 2: Key influencer -
A key influencer chart displays the major contributors to a selected result or value.

Box 3: Scatter -
Scatter and Bubble charts display relationships between 2 (scatter) or 3 (bubble) quantitative measures --whether or not, in which order, etc.

Question: 16

You need to create an Azure Storage account.
Data in the account must replicate outside the Azure region automatically.
Which two types of replication can you use for the storage account? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. zone-redundant storage (ZRS)
B. read-access geo-redundant storage (RA-GRS)
C. locally-redundant storage (LRS)
D. geo-redundant storage (GRS)

Answer: BD

Explanation:
D: Azure Storage offers two options for copying your data to a secondary region:
☞ Geo-redundant storage (GRS)
☞ Geo-zone-redundant storage (GZRS)
B: With GRS or GZRS, the data in the secondary region isn't available for read or write access unless there is a failover to the secondary region. For read access to the secondary region, configure your storage account to use read-access geo-redundant storage (RA-GRS) or read-access geo-zone-redundant storage (RA-GZRS).

Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy#redundancy-in-a-secondary-r egion

Question: 17

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:

Answer Area

Statements	Yes	No
Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than infrastructure as a service (IaaS) database offerings.	○	○
Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.	○	○
All relation and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.	○	○

Answer:

Answer Area

Statements	Yes	No
Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than infrastructure as a service (IaaS) database offerings.	○	○
Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.	○	○
All relation and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.	○	○

Explanation:
Box 1: Yes -
Like IaaS, PaaS includes infrastructure " servers, storage, and networking " but also middleware, development tools, business intelligence (BI) services, database management systems, and more. PaaS is designed to support the complete web application lifecycle: building, testing, deploying, managing, and updating.

PaaS allows you to avoid the expense and complexity of buying and managing software licenses, the underlying application infrastructure and middleware, container orchestrators such as Kubernetes, or the development tools and other resources

Box 2: No -
You manage the applications and services you develop, and the cloud service provider typically manages everything else.

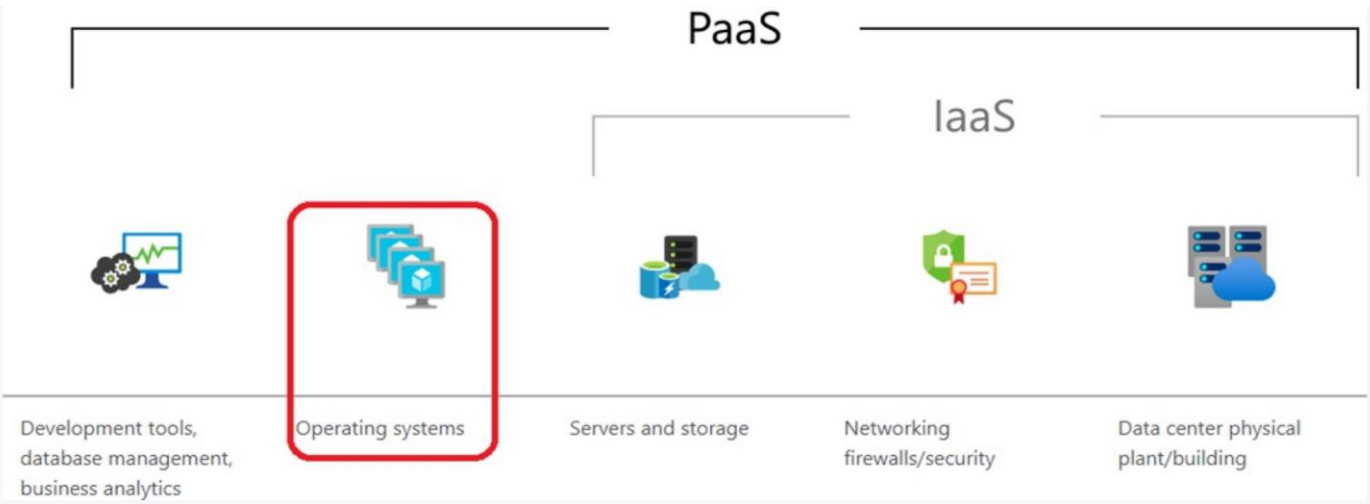

Box 3: No -
There really is no way to pause / stop billing for your Azure SQL Database.

Microsoft's official answer "Yes, you can export your database. Delete the Azure SQL database and that will pause billing. Then when you need it you can create a new database and import your previously exported DB."

Reference:
https://azure.microsoft.com/en-us/overview/what-is-paas

Question: 18

Which statement is an example of Data Manipulation Language (DML)?

A. REVOKE
B. DISABLE
C. INSERT
D. GRANT

Answer: C

Explanation:
Data Manipulation Language (DML) affect the information stored in the database. Use these statements to insert, update, and change the rows in the database.

BULK INSERT -

- DELETE
- INSERT
- SELECT
- UPDATE
- MERGE

Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/statements

Question: 19

You have a SQL query that combines customer data and order data. The query includes calculated columns. You need to create a database object that would allow other users to rerun the same SQL query.
What should you create?

A. an index
B. a view
C. a scalar function
D. a table

Answer: B

Explanation:
A view is a virtual table whose contents are defined by a query. A view acts as a filter on the underlying tables referenced in the view. The query that defines the view can be from one or more tables or from other views in the current or other databases.

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

Question: 20

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

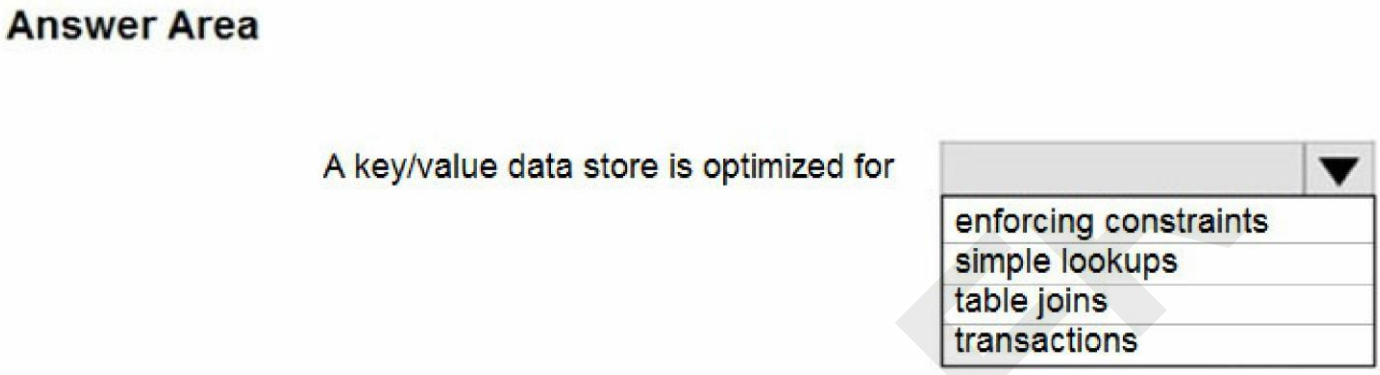

Answer:

Explanation:
Box 1: simple lookups -
A key/value store associates each data value with a unique key. Most key/value stores only support simple query, insert, and delete operations. To modify a value
(either partially or completely), an application must overwrite the existing data for the entire value. In most implementations, reading or writing a single value is an atomic operation.

An application can store arbitrary data as a set of values. Any schema information must be provided by the application. The key/value store simply retrieves or stores the value by key.

Reference:
https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview

Question: 21 DRAG

DROP -

Match the types of data to the appropriate Azure data services.

To answer, drag the appropriate data type from the column on the left to its service on the right. Each data type may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:

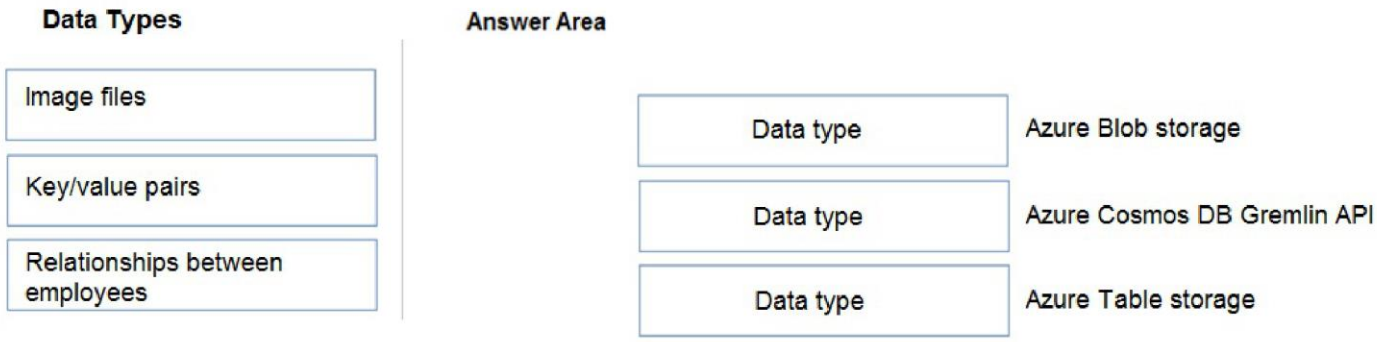

Answer:

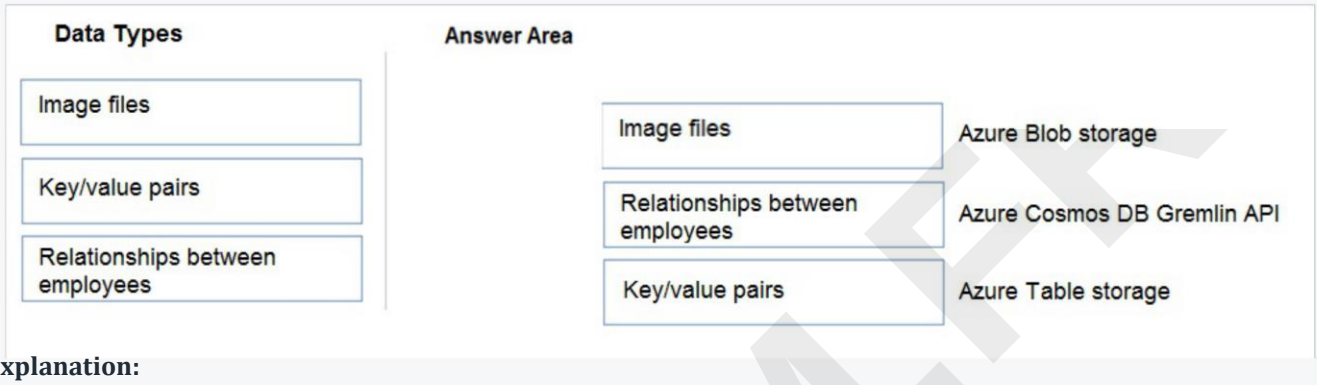

Explanation:

1. Image files = Azure Blob storage

2. Relationship between employyes = Azure Cosmos DB Gremlin API

3. Key/value pairs = Azure Table Storage

Azure Blob Storage is designed for storing large amounts of unstructured data, such as images, videos, and documents.

Azure Table Storage is a NoSQL store that efficiently handles structured data in the form of key/value pairs. Azure Cosmos DB Gremlin API is used for graph-based data, which is ideal for representing relationships, such as the hierarchy or connections between employees.

Reference:
https://docs.microsoft.com/en-us/azure/storage/tables/table-storage-design-modeling

Question: 22

DRAG DROP -
Match the Azure Data Lake Storage Gen2 terms to the appropriate levels in the hierarchy.
To answer, drag the appropriate term from the column on the left to its level on the right. Each term may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.
Select and Place:

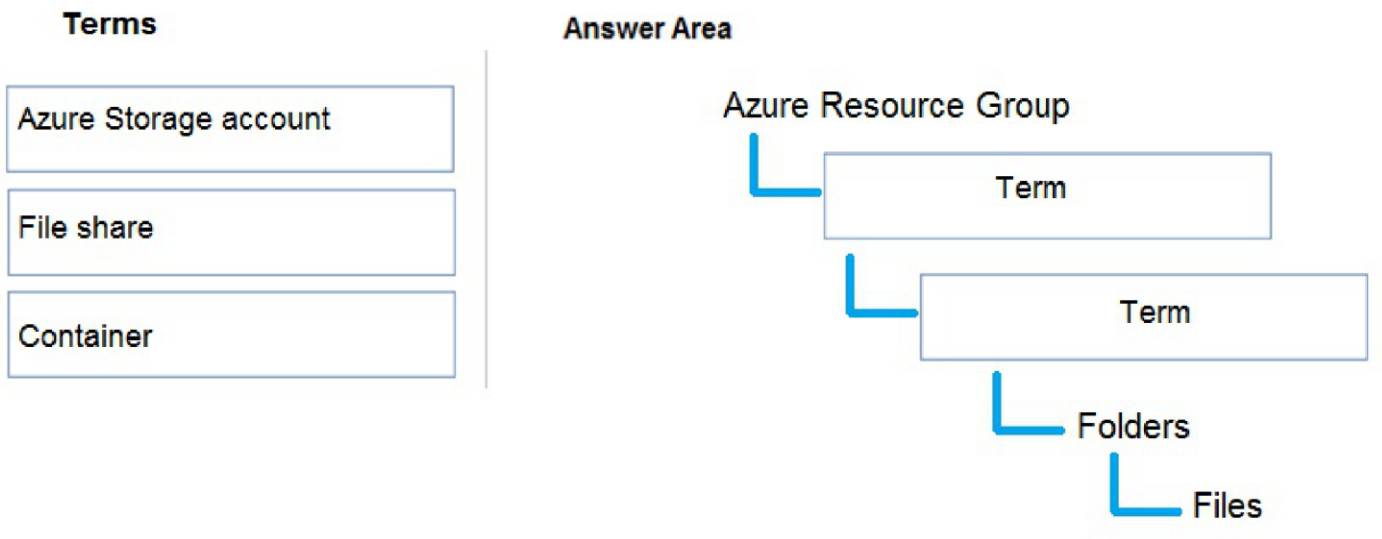

Answer:

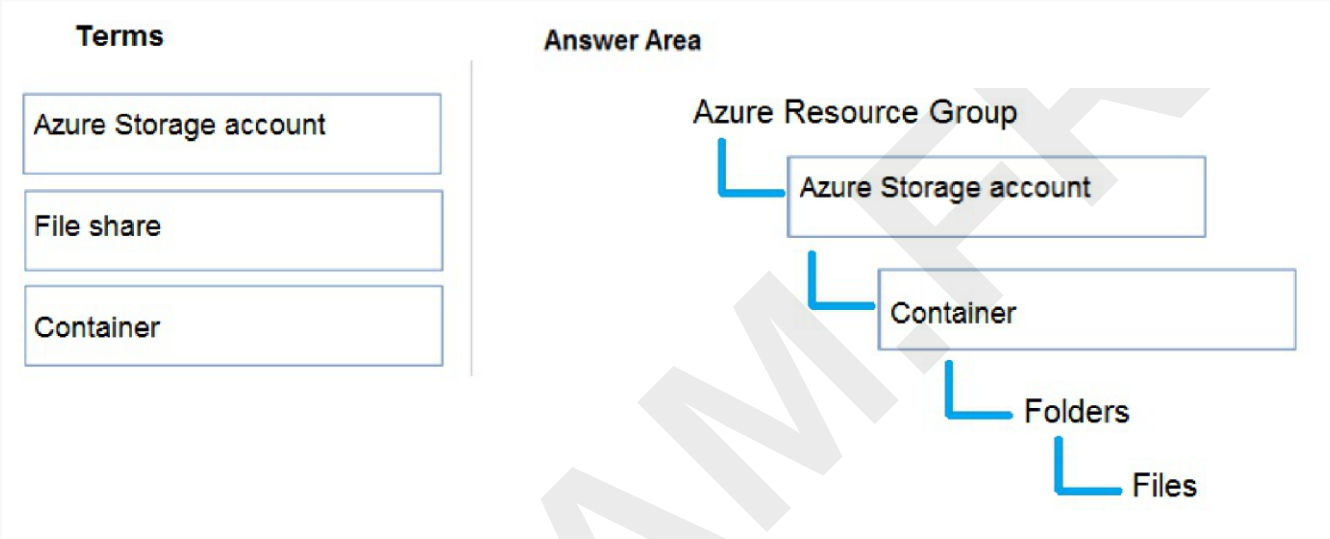

Explanation:

Box 1: Azure Storage account -

Azure file shares are deployed into storage accounts, which are top-level objects that represent a shared pool of storage.

Box 2: Container -

Azure Storage Account >> Container >> Directories >> Files

Azure Data Lake Storage Gen2 builds on blob storage and optimizes I/O of high-volume data by using a hierarchical namespace that organizes blob data into directories, and stores metadata about each directory and the files within it. This structure allows operations, such as directory renames and deletes, to be

performed in a single atomic operation. Flat namespaces, by contrast, require several operations proportionate to the number of objects in the structure. Hierarchical namespaces keep the data organized, which yields better storage and retrieval performance for an analytical use case and lowers the cost of analysis.

SOURCE: https://learn.microsoft.com/en-us/training/modules/introduction-to-azure-data-lake-storage/4-azure-data-lake-and-blob-storage

Reference:

https://docs.microsoft.com/en-us/azure/storage/files/storage-how-to-create-file-share

Question: 23

What are two characteristics of real-time data processing? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. Data is processed periodically
B. Low latency is expected
C. High latency is acceptable
D. Data is processed as it is created

Answer: BD

Explanation:
Real time processing deals with streams of data that are captured in real-time and processed with minimal latency to generate real-time (or near-real-time) reports or automated responses.

Reference:
https://docs.microsoft.com/en-us/azure/architecture/data-guide/big-data/real-time-processing

Question: 24

DRAG DROP -
Match the Azure Data Factory components to the appropriate descriptions.
To answer, drag the appropriate component from the column on the left to its description on the right. Each component may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.
Select and Place:

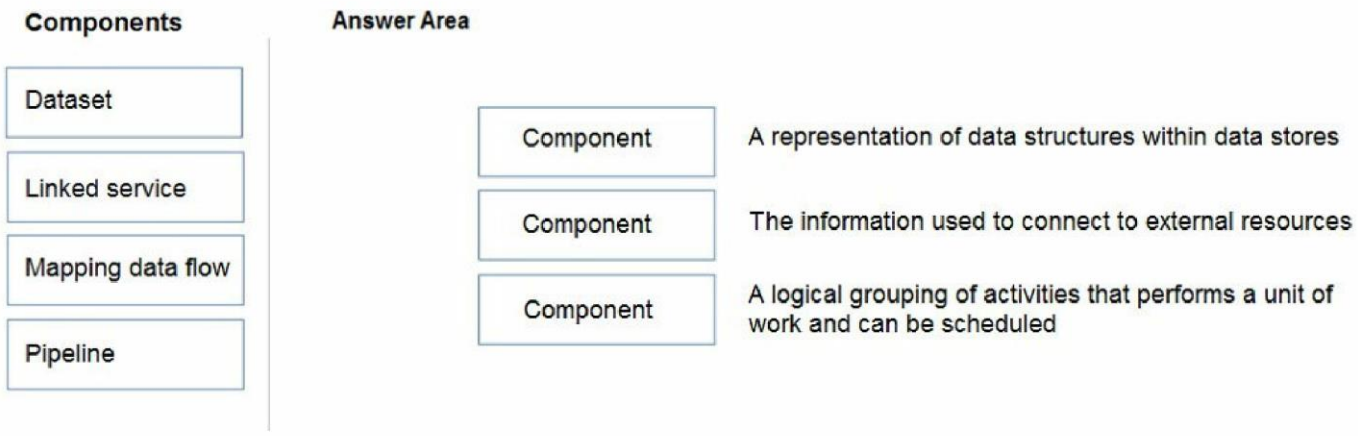

Answer:

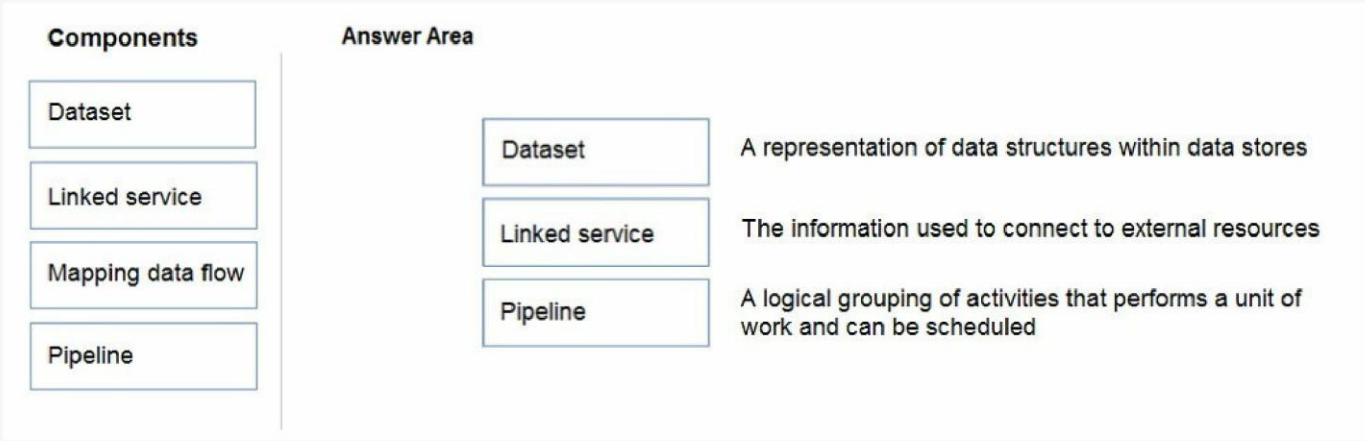

Explanation:
Box 1: Dataset -
Datasets must be created from paths in Azure datastores or public web URLs, for the data to be accessible by Azure Machine Learning.

Box 2: Linked service -
Linked services are much like connection strings, which define the connection information needed for Data Factory to connect to external resources.

Box 3: Pipeline -
A pipeline is a logical grouping of activities that together perform a task.

Reference:
https://k21academy.com/microsoft-azure/dp-100/datastores-and-datasets-in-azure/ https://docs.microsoft.co m/en-us/azure/data-factory/concepts-linked-services https://docs.microsoft.com/en-us/azure/data-factory/c oncepts-pipelines-activities

Question: 25

DRAG DROP -
Match the types of workloads to the appropriate scenarios.
To answer, drag the appropriate workload type from the column on the left to its scenario on the right. Each workload type may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.
Select and Place:

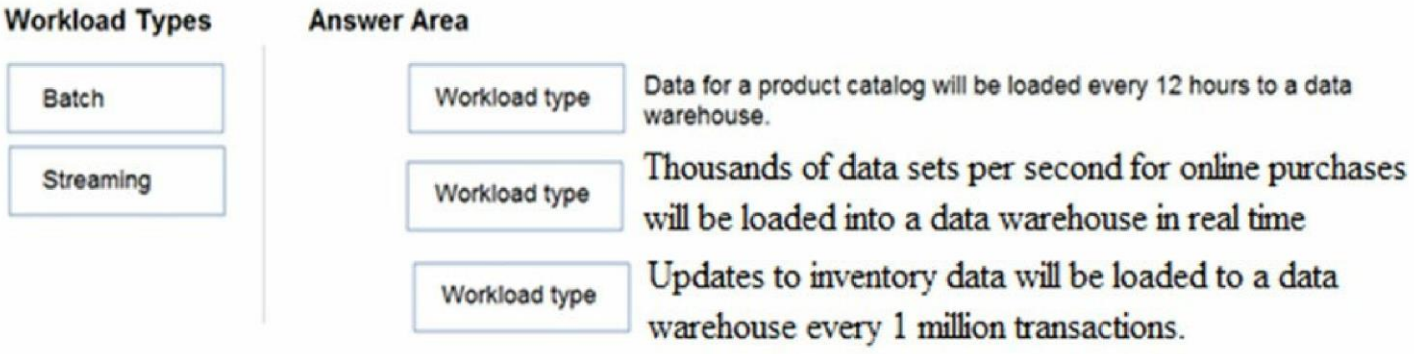

Answer:

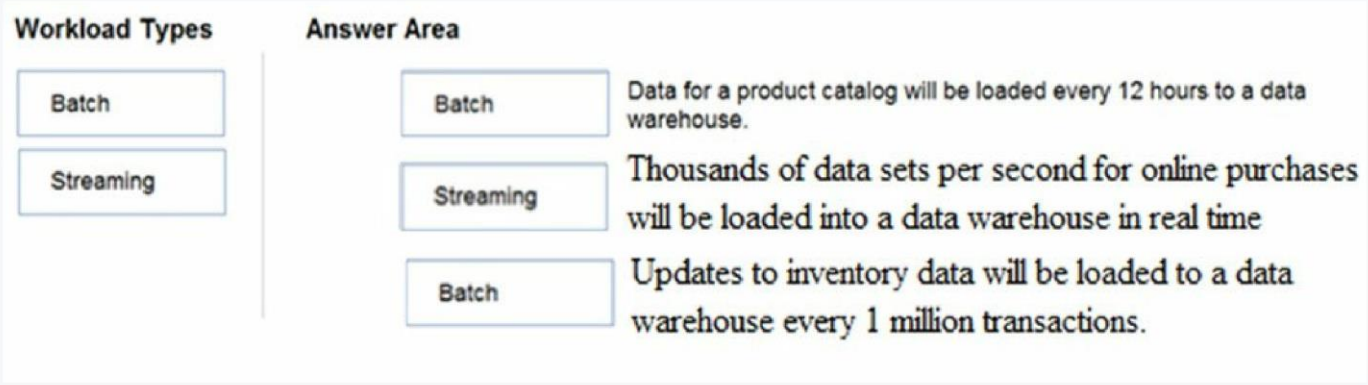

Explanation:
Box 1: Batch -
Batch processing refers to the processing of blocks of data that have already been stored over a period of time.

Box 2: Streaming -
Stream processing is a big data technology that allows us to process data in real-time as they arrive and detect conditions within a small period of time from the point of receiving the data. It allows us to feed data into analytics tools as soon as they get generated and get instant analytics results.

Box 3: Batch -

Reference:
https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/batch-processing

Question: 26

DRAG DROP -
Your company plans to load data from a customer relationship management (CRM) system to a data warehouse by using an extract, load, and transform (ELT) process.
Where does data processing occur for each stage of the ELT process? To answer, drag the appropriate locations to the correct stages. Each location may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.
Select and Place:

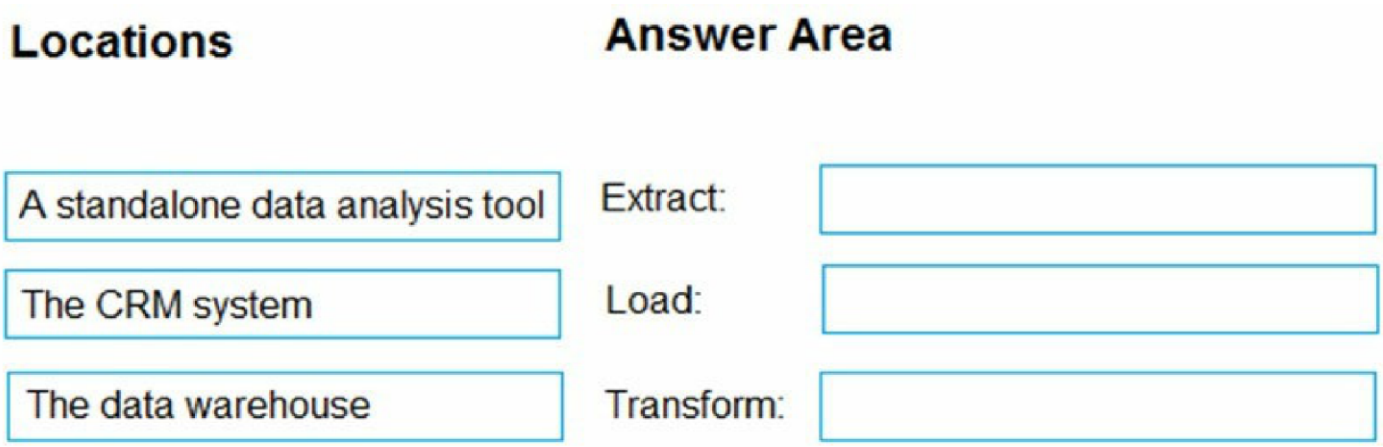

Answer:

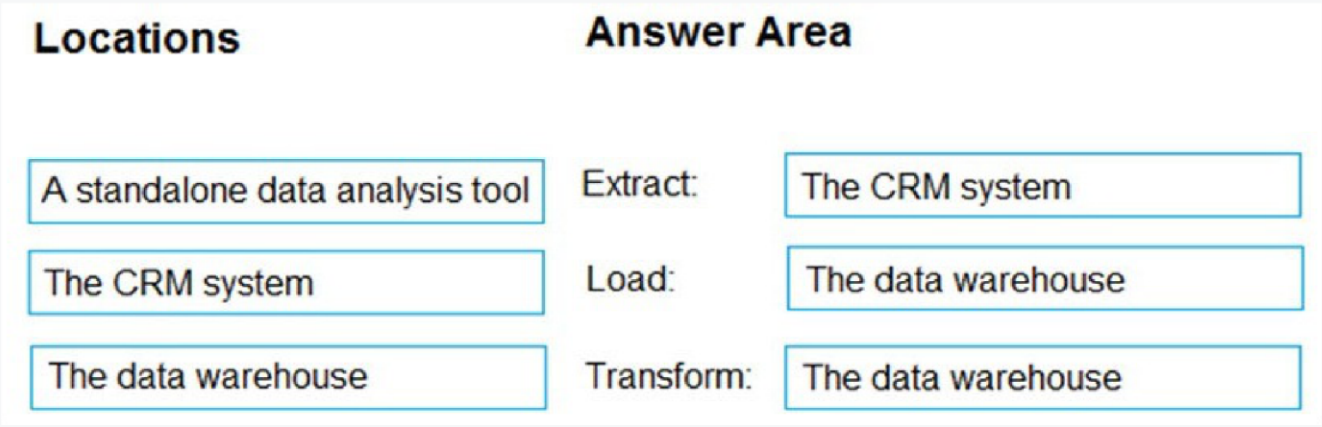

Explanation:

Box 1: The CRM system -

Data is extracted from the CRM system.

Box 2: The data warehouse -

Data is loaded to the data warehouse.

Box 3: The data warehouse -

The data warehouse, In an ELT process, the transformation of data occurs within the target system (data warehouse), leveraging its processing power for transformations.

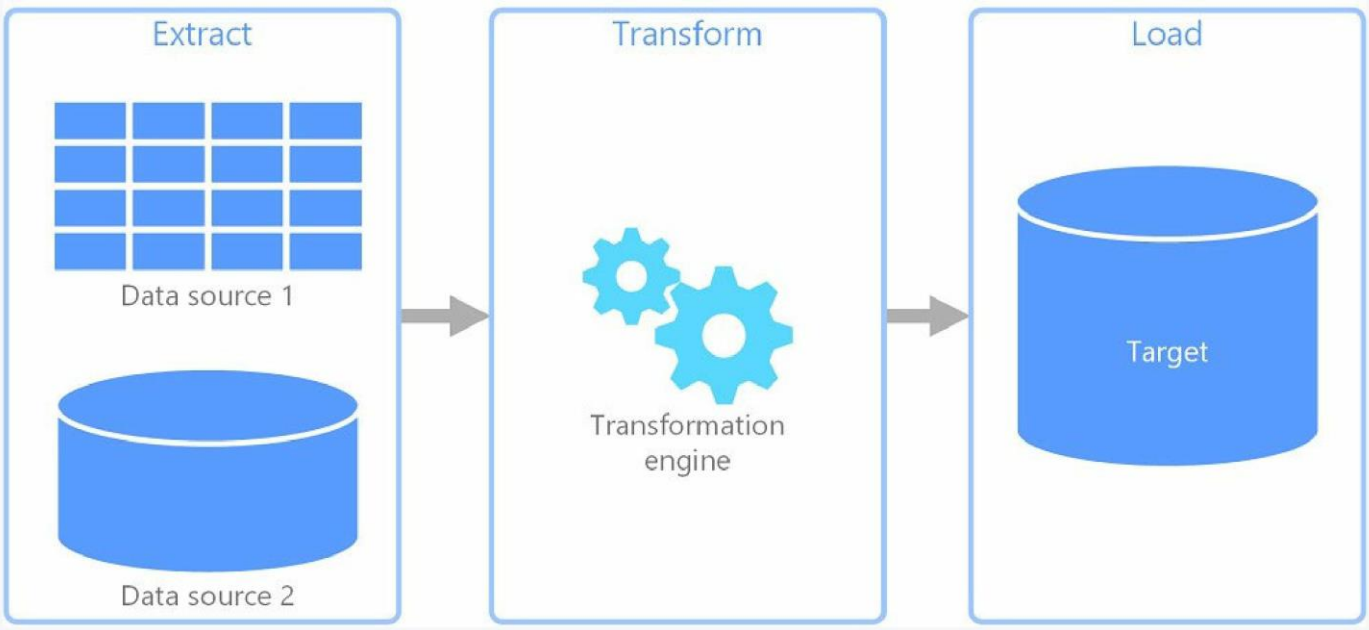

Reference:

https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl

Question: 27

Your company recently reported sales from the third quarter. You have the chart shown in the following exhibit.

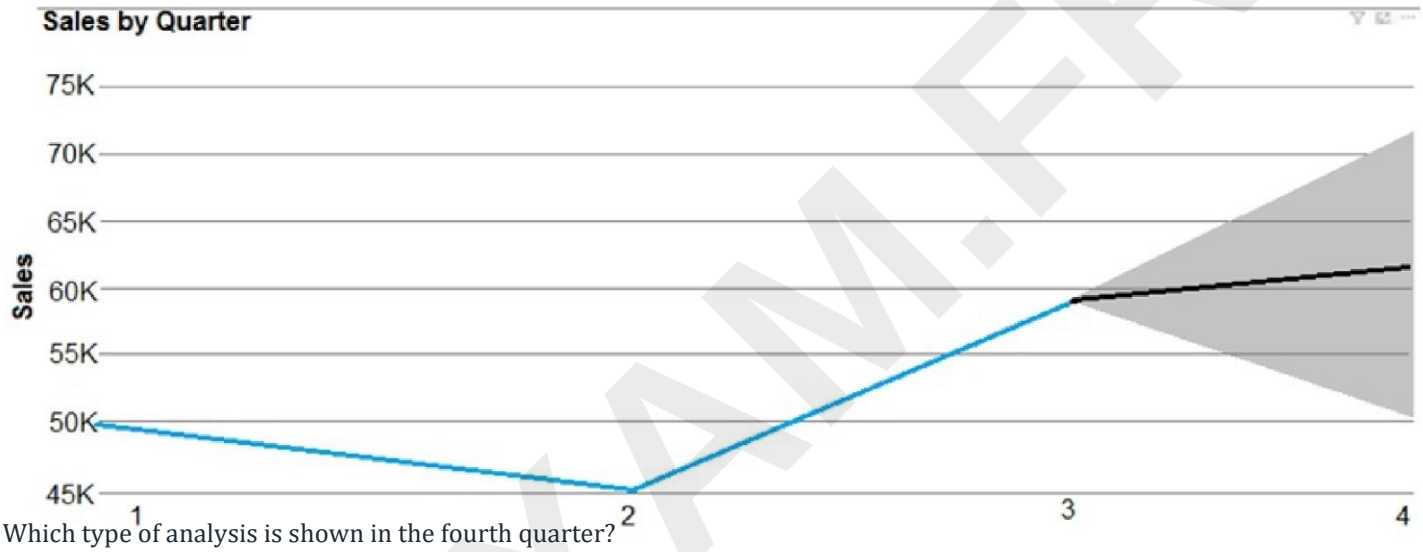

Which type of analysis is shown in the fourth quarter?

A. predictive
B. prescriptive
C. descriptive
D. diagnostic

Answer: A

Explanation:
Predictive, to answer the question: What will happen?

Reference:
https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cogniti ve-analytics/ https://azure.microsoft.com/en-us/blog/answering-whats-happening-whys-happening-and-what-will-happen-with-iot-analytics/

Question: 28 Which statement is an example of Data Manipulation Language (DML)?

A. REVOKE
B. DISABLE
C. CREATE
D. UPDATE

Answer: D

Explanation:
Data Manipulation Language (DML) affect the information stored in the database. Use these statements to insert, update, and change the rows in the database.

BULK INSERT -

DELETE -

INSERT -

SELECT -

UPDATE -

MERGE -

There's 3 type of Relational data SQL:

DDL (Data Definition Language) CREATE, DROP, ALTER, RENAME, COMMENT AND TRUNCATE.

DML (Data Manipulación Language) SELECT, INSERT INTO, DELETE, UPDATE.

DCL (Data Control Language) REVOKE & GRANT.

TCL (Transaccional Control Language) COMMIT & ROLLBACK

Reference:

https://docs.microsoft.com/en-us/sql/t-sql/statements/statements

Question: 29 HOTSPOT

-

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

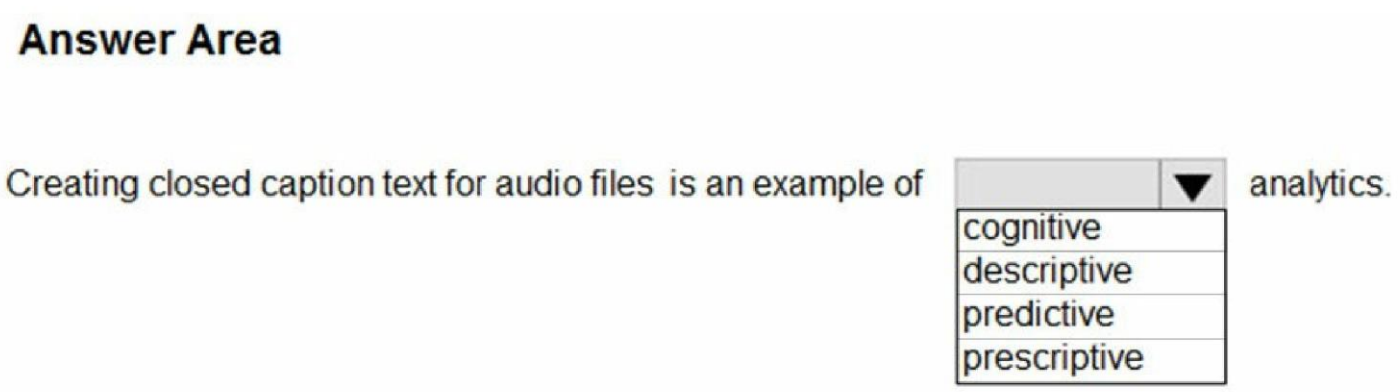

Answer:

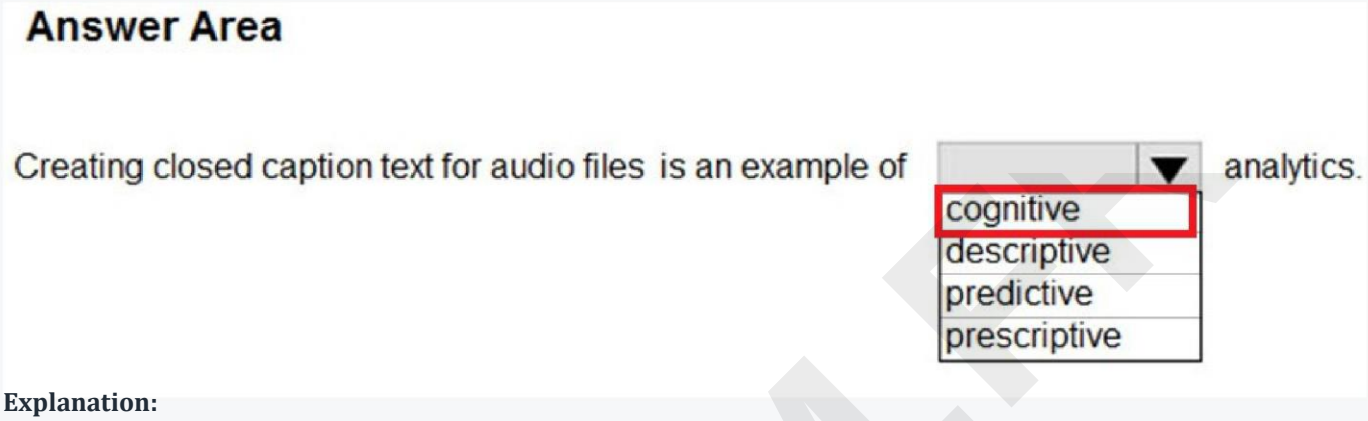

Explanation:

Correct answer is Cognitive.

Reference:

https://azure.microsoft.com/en-us/services/cognitive-services/speech-to-text/#overview

Question: 30 HOTSPOT

-

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

Answer Area

A visualization that illustrates a university's current student enrollment per department is an example of [▼] analytics.

- cognitive
- descriptive
- predictive
- prescriptive

Answer:

Answer Area

A visualization that illustrates a university's current student enrollment per department is an example of [▼] analytics.

- cognitive
- descriptive
- predictive
- prescriptive

Explanation:
Generally speaking, data analytics comes in four types:
1. Descriptive, to answer the question: What's happening?
2. Diagnostic, to answer the question: Why's happening?
3. Predictive, to answer the question: What will happen?
4. Prescriptive, to answer the question: What actions should we take?

Reference:
https://azure.microsoft.com/en-us/blog/answering-whats-happening-whys-happening-and-what-will-happen-with-iot-analytics/

Question: 31

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

Answer Area

An extract, load, and transform (ELT) process requires [▼]

- a separate transformation engine.
- a target data store powerful enough to transform data.
- data that is fully processed before being loaded to the target data store.
- a data pipeline that includes a transformation engine.

Answer:

Answer Area

An extract, load, and transform (ELT) process requires [▼]

- a separate transformation engine.
- a target data store powerful enough to transform data.
- data that is fully processed before being loaded to the target data store.
- a data pipeline that includes a transformation engine.

Explanation:
With ELT, the data store used to perform the transformation is the same data store where the data is ultimately consumed.

Reference:
https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl

Question: 32

You need to create an Azure resource to store data in Azure Table storage. Which command should you run?

A. az storage share create
B. az storage account create
C. az cosmosdb create
D. az storage container create

Answer: B

Explanation:

Due to no az storage table create in the selection. The best answer is B create storage account (first)

Question: 33

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

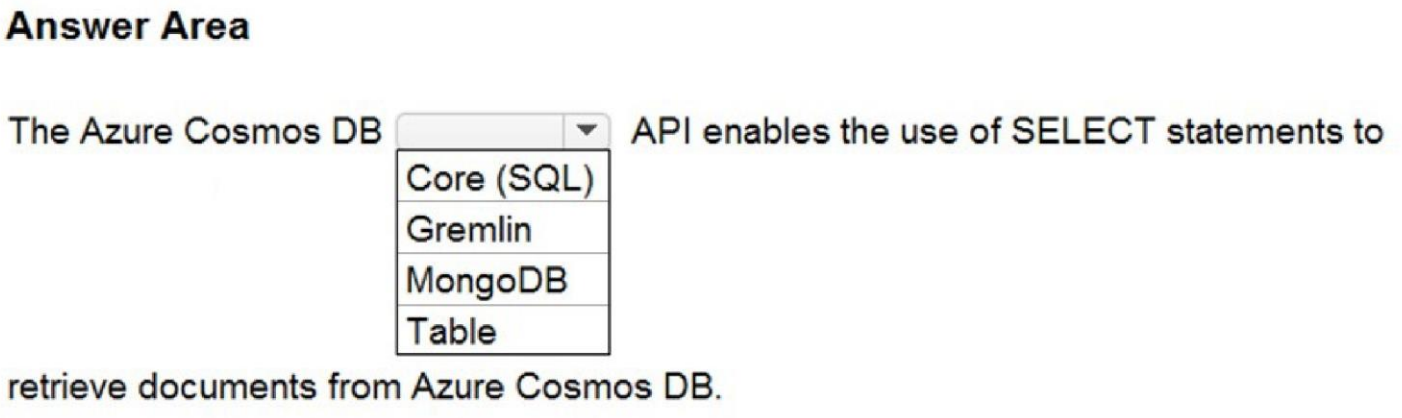

Answer:

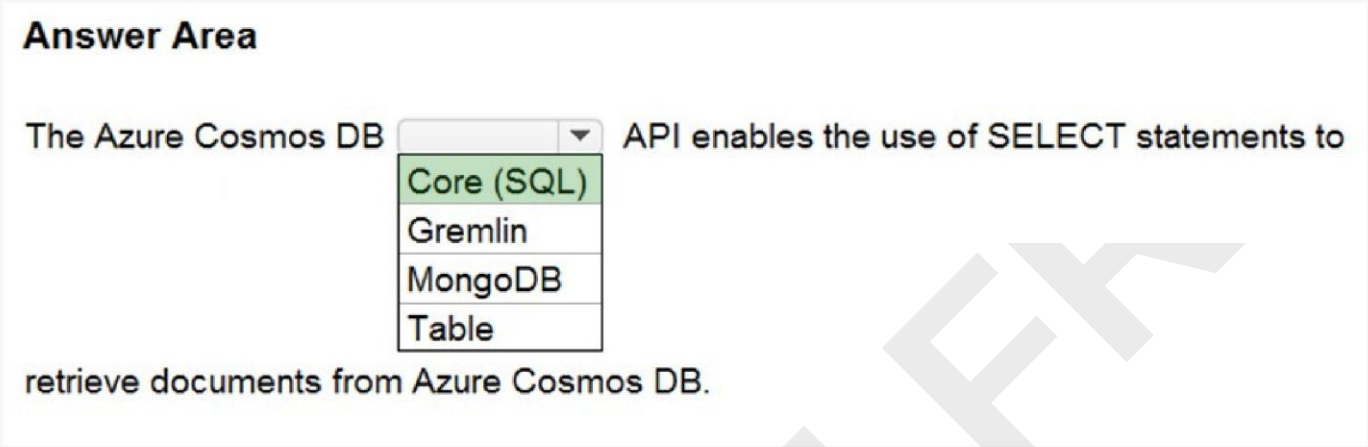

Explanation:
Azure Cosmos DB SQL API accounts provide support for querying items using the Structured Query Language (SQL) syntax.

Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api

Question: 34

You need to modify a view in a relational database by adding a new column. Which statement should you use?

A. MERGE
B. ALTER
C. INSERT
D. UPDATE

Answer: B

Explanation:

"Modify a View" = Alter View ....

There's 3 type of Relational data SQL:

DDL (Data Definition Language) CREATE, DROP, ALTER, RENAME, COMMENT AND TRUNCATE. DML (Data Manipulación Language) SELECT, INSERT INTO, DELETE, UPDATE.

DCL (Data Control Language) REVOKE & GRANT.

TCL (Transactional Control Language) COMMIT & ROLLBACK

Question: 35 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:

Answer Area

Statements	Yes	No
Extract, transform, and load (ETL) can reduce the transfer of sensitive data to destination systems.	◯	◯
Extract, load, and transform (ELT) transforms data by using a compute resource independent of the source system and destination system.	◯	◯
Extract, load, and transform (ELT) minimizes the time it takes to copy large volumes of data to destination systems.	◯	◯

Answer:

Answer Area

Statements	Yes	No
Extract, transform, and load (ETL) can reduce the transfer of sensitive data to destination systems.	◯	◯
Extract, load, and transform (ELT) transforms data by using a compute resource independent of the source system and destination system.	◯	◯
Extract, load, and transform (ELT) minimizes the time it takes to copy large volumes of data to destination systems.	◯	◯

Explanation:

"Extract, transform, and load (ETL) can reduce the transfer of sensitive data to destination systems." → **Yes** ETL processes data before loading it into the destination system. This means that sensitive or unnecessary data can be filtered out during the transformation step, reducing exposure and improving security.

This helps in compliance and data governance by ensuring that only relevant and sanitized data reaches the destination system.

"Extract, load, and transform (ELT) transforms data by using a compute resource independent of the source system and destination system." →**No**
In ELT, data is first loaded into the destination system (usually a data warehouse like Snowflake, BigQuery, or Redshift), and then transformations occur within the destination system itself.

This means that ELT relies on the compute power of the destination system, not an independent compute resource.

Unlike ETL, which transforms data before loading, ELT performs transformations after loading, leveraging the destination system's resources.

"Extract, load, and transform (ELT) minimizes the time it takes to copy large volumes of data to destination systems." →**Yes**
ELT is optimized for handling large datasets because it loads data in bulk first and then processes it within the destination system.

This approach is faster than ETL for large-scale data processing because data can be loaded quickly without transformation overhead, and modern cloud-based data warehouses provide powerful compute resources for transformation.

It is commonly used in big data and cloud-based analytics environments where storage and compute are scalable.

Question: 36

HOTSPOT -
You plan to deploy a PostgreSQL database to Azure.
Which hosting model corresponds to the available deployment options? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

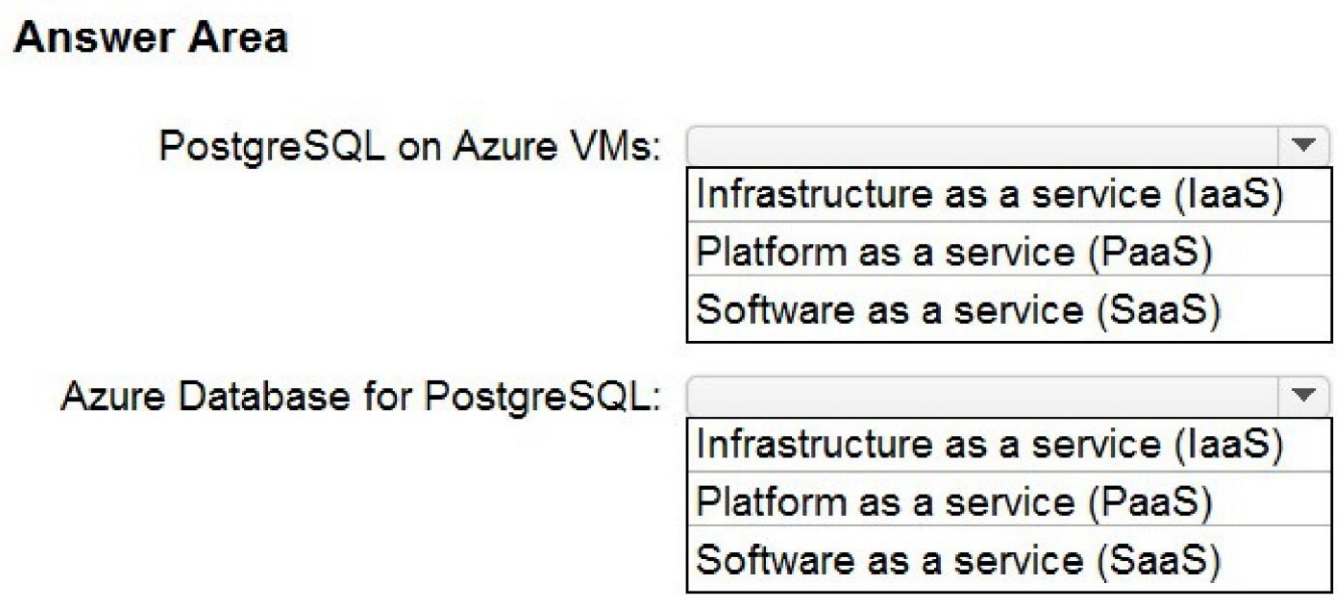

Answer:

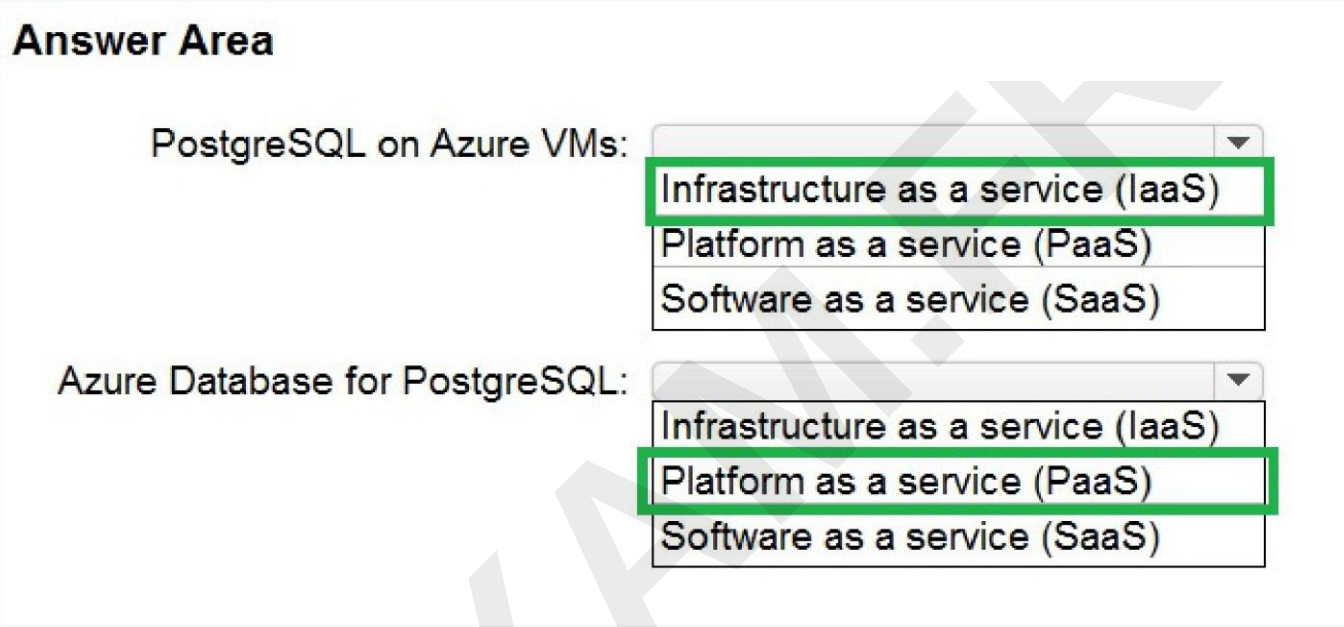

Explanation:

PostgreSQL on Azure VMs

1. IaaS (Infrastructure as a Service)

Deploying PostgreSQL on an Azure Virtual Machine means you are responsible for managing the virtualized infrastructure, including the operating system, database installation, updates, and backups. This falls under the IaaS model.

Azure Database for PostgreSQL

2. PaaS (Platform as a Service)
Azure Database for PostgreSQL is a fully managed database service where Azure handles the infrastructure, scaling, updates, backups, and security. You only manage the database and data, making it a PaaS solution.

Reference:

https://azure.microsoft.com/en-us/overview/what-is-saas/

https://azure.microsoft.com/en-us/overview/what-is-paas/

Question: 37 HOTSPOT

-

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

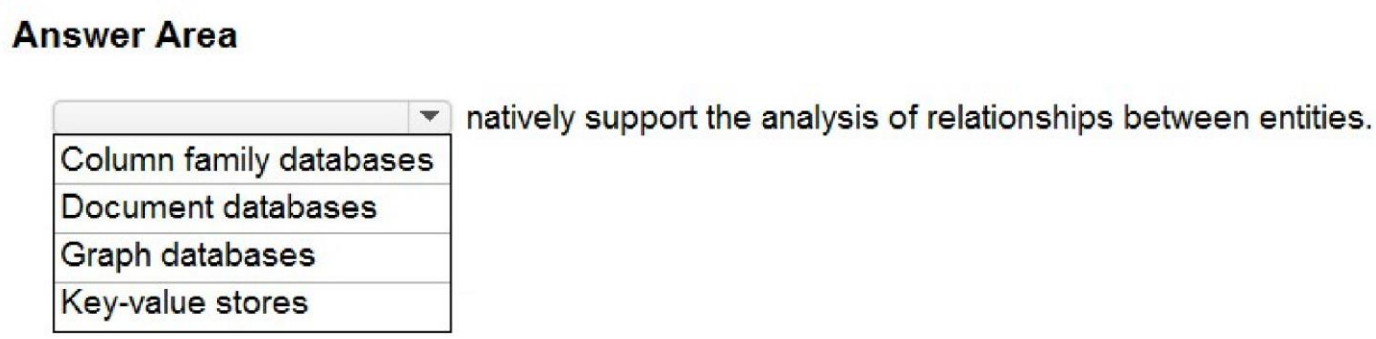

Answer:

Explanation:

Graph databases:

Designed specifically to represent and query relationships between entities using nodes and edges. They are optimized for queries involving complex relationships, such as social networks or recommendation engines.

Why the other options are incorrect:

Column family databases: These store data in column families (e.g., Cassandra). They are optimized for wide-row queries but do not inherently support relationship analysis.

Document databases: Store data as documents (e.g., MongoDB). While they allow for nested structures, they

do not specialize in managing or querying relationships.

Key-value stores: Store data as simple key-value pairs (e.g., Redis). They are ideal for quick lookups but lack native features for relationship analysis.

Important Tip:

When analyzing relationships between entities, always consider Graph databases (e.g., Neo4j) for their efficiency and built-in support for relationship-based queries.

Question: 38

Which Azure storage solution provides native support for POSIX-compliant access control lists (ACLs)?

A. Azure Table storage
B. Azure Data Lake Storage
C. Azure Queue storage
D. Azure Files

Answer: B

Explanation:

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

Question: 39

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

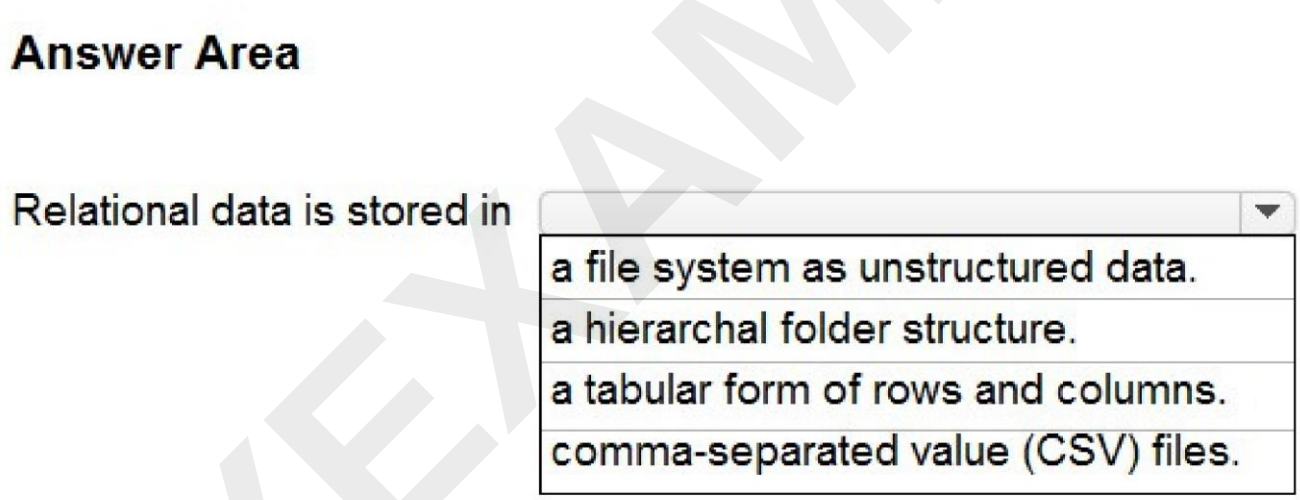

Answer:

Answer Area

Relational data is stored in

- a file system as unstructured data.
- a hierarchal folder structure.
- a tabular form of rows and columns.
- comma-separated value (CSV) files.

Question: 40 HOTSPOT

-

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

Answer Area

- A view
- A clustered index
- A stored procedure
- A nonclustered index

physically sorts the data in a table based on the values in a specified column.

Answer:

Answer Area

- A view
- A clustered index
- A stored procedure
- A nonclustered index

physically sorts the data in a table based on the values in a specified column.

Explanation:

Answer is B, Clustered indexes.

https://docs.microsoft.com/en-us/sql/relational-databases/indexes/clustered-and-nonclustered-indexes-described?view=sql-server-ver15&viewFallbackFrom=sql-server-ver15%20%20%20Previous%20QuestionsNext%20Questions

Reference:

https://docs.microsoft.com/en-us/sql/relational-databases/indexes/clustered-and-nonclustered-indexes-described?view=sql-server-ver15

Question: 41

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

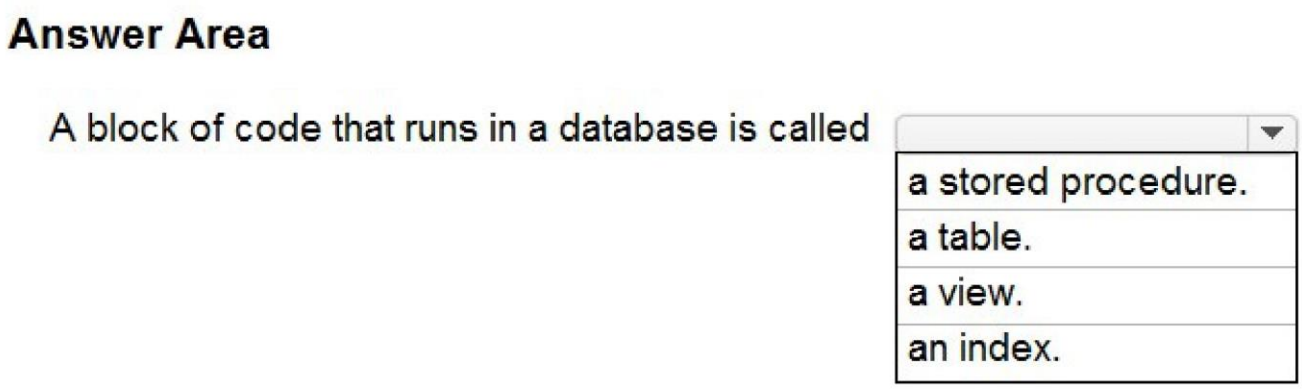

Answer:

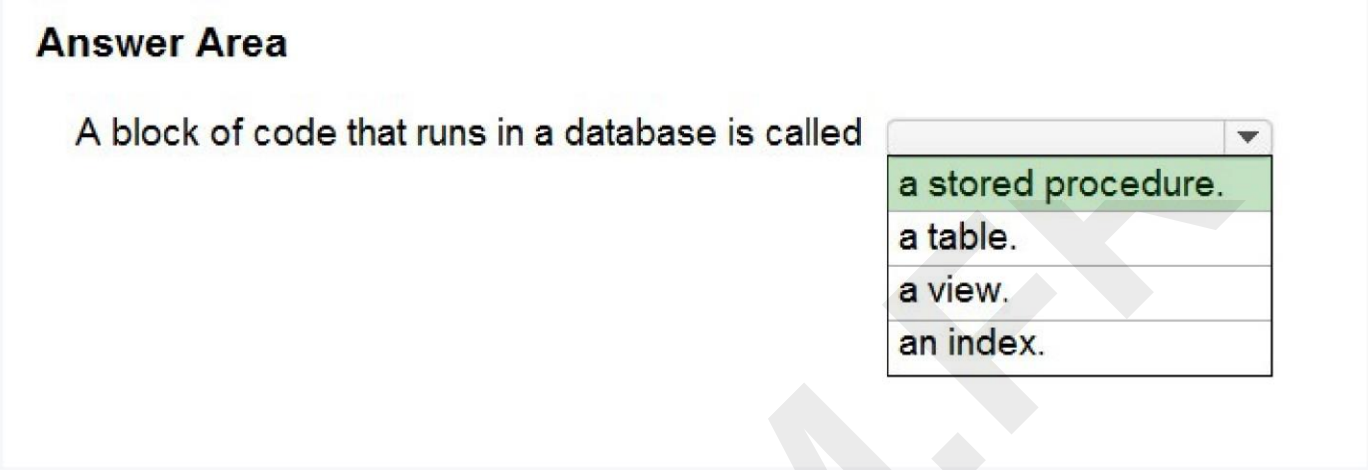

Explanation:

Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/stored-procedures/stored-procedures-database-engine?view=sql-server-ver15

Question: 42

DRAG DROP -
Match the types of analytics that can be used to answer the business questions.
To answer, drag the appropriate analytics type from the column on the left to its question on the right. Each analytics type may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.
Select and Place:

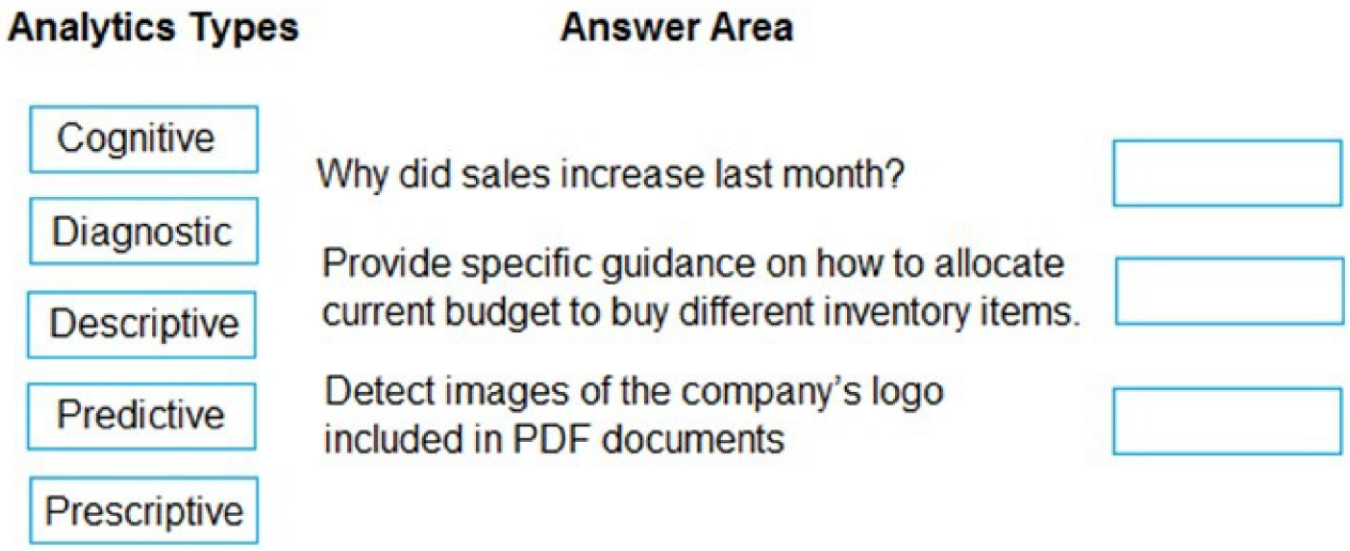

Answer:

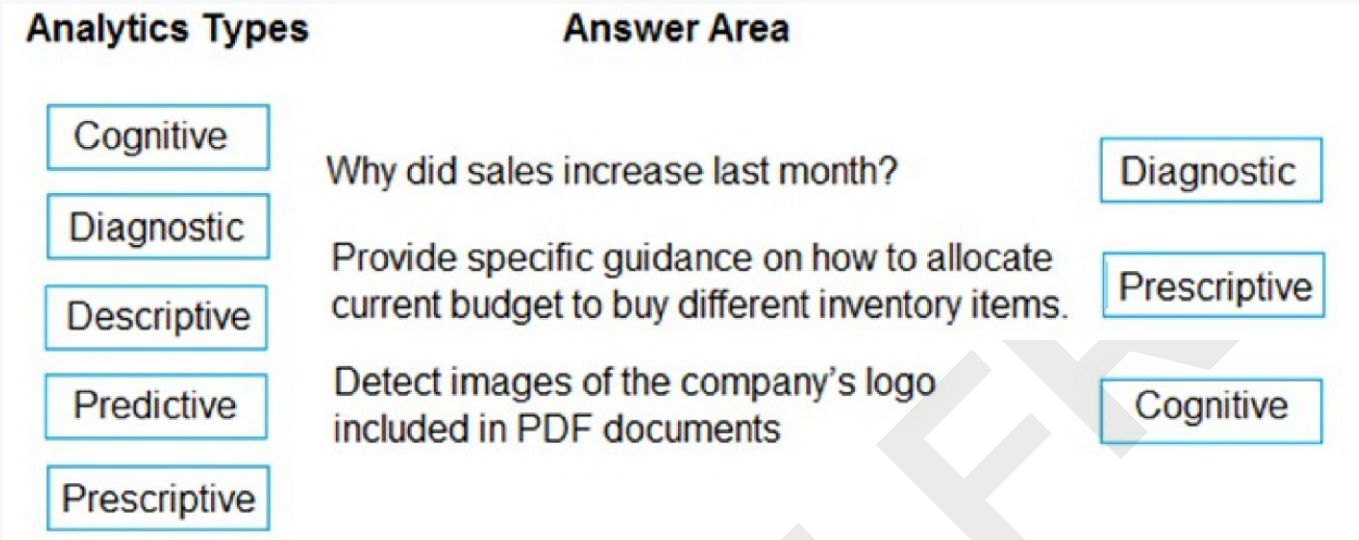

Explanation:

Box 1: Diagnostic -

Diagnostic Analytics: At this stage you can begin to answer some of those why questions. Historical data can begin to be measured against other data to answer the question of why something happened in the past. This is the process of gathering and interpreting different data sets to identify anomalies, detect patters, and determine relationships.

Box 2: Prescriptive -

Prescriptive analytics is a combination of data, mathematical models, and various business rules to infer actions to influence future desired outcomes.

Predictive analytics, broadly speaking, is a category of business intelligence that uses descriptive and predictive variables from the past to analyze and identify the likelihood of an unknown future outcome Box 3: Cognitive -

Reference:

https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitive-analytics/ https://azure.microsoft.com/en-us/blog/answering-whats-happening-whys-happening-

and-what-will-happen-with-iot-analytics/

Question: 43

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

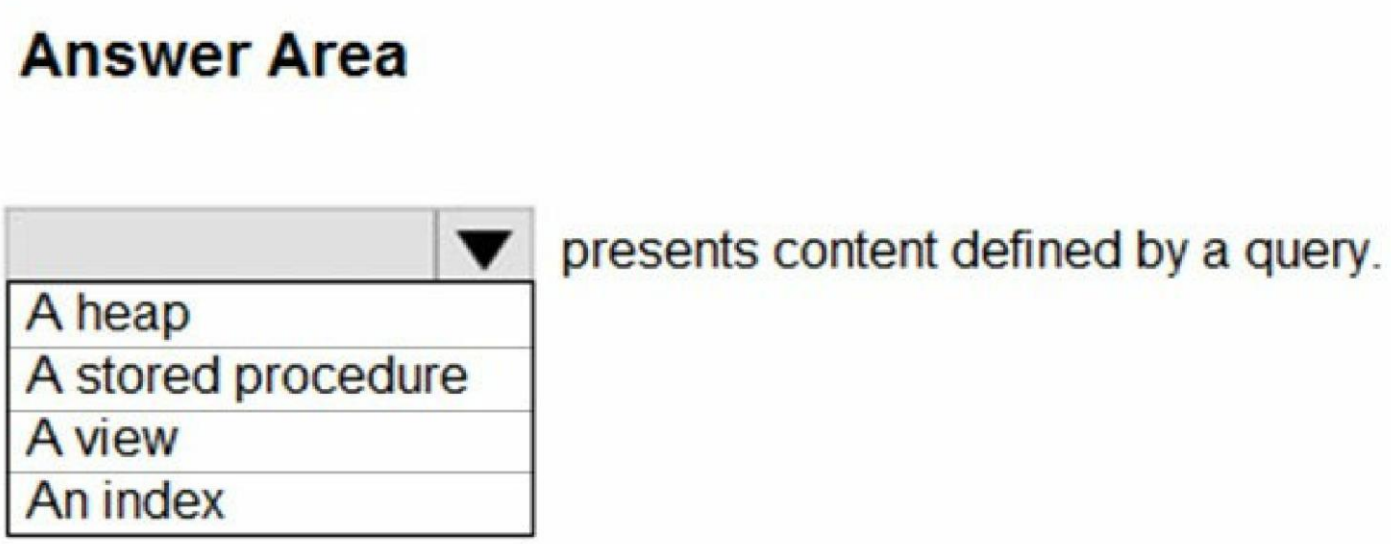

Answer:

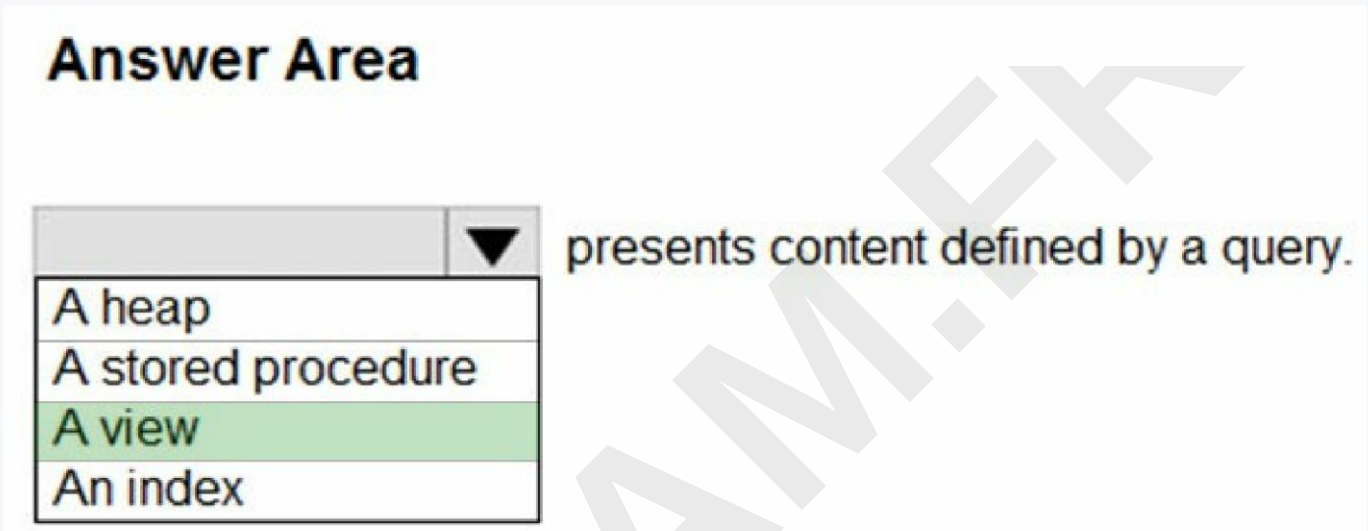

Question: 44

Which type of database is Azure Database for PostgreSQL?

A. Platform as a service (PaaS)
B. Infrastructure as a service (IaaS)
C. Microsoft SQL Server
D. on-premises

Answer: A

Explanation:

Correct answer: A, PAAS

Reference:

https://docs.microsoft.com/en-us/azure/postgresql/overview-postgres-choose-server-options

Question: 45

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

Answer: A

Explanation:
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: 46

HOTSPOT -
Select the answer that correctly completes the sentence.
Hot Area:

Answer Area

A [▼] is a chart of colored rectangles. The size of the rectangles represent the relative value of each item. The chart can be hierarchical, displaying data as a set of rectangles nested within the main rectangle.

- line chart
- matrix
- scatter
- treemap

Answer:

Answer Area

A [▼] is a chart of colored rectangles. The size of the rectangles represent the relative value of each item. The chart can be hierarchical, displaying data as a set of rectangles nested within the main rectangle.

- line chart
- matrix
- scatter
- treemap

Explanation:
A treemap chart divides the chart area into rectangles that represent the different levels and relative sizes of the data hierarchy.

Reference:
https://docs.microsoft.com/en-us/sql/reporting-services/report-design/tree-map-and-sunburst-charts-in-repo rting-services?view=sql-server-ver15

Question: 47

What is a characteristic of batch processing?

A. The data ingested during batch processing must be processed as soon as the data is received.
B. Large datasets must be split into batches of less than 1 GB before the data can be processed.
C. There is a noticeable time delay between ingesting data and obtaining the data processing results.
D. Batch processing can only process data that is structured.

Answer: C

Explanation:

C. Batch processing can process both structured and unstructured data.

Question: 48

HOTSPOT -
Select the answer that correctly completes the sentence.
Hot Area:

Answer Area

In a SQL statement, [▼] are examples of aggregate functions.

- FROM and WHERE
- GROUP BY and ORDER BY
- JOIN and MERGE
- MAX and MIN

Answer:

Answer Area

In a SQL statement, [▼] are examples of aggregate functions.

- FROM and WHERE
- GROUP BY and ORDER BY
- JOIN and MERGE
- MAX and MIN

Explanation:
An aggregate function performs a calculation on a set of values, and returns a single value. The following are aggregate functions:

- APPROX_COUNT_DISTINCT
- AVG
- CHECKSUM_AGG
- COUNT
- COUNT_BIG

GROUPING -

- GROUPING_ID
- MAX
- MIN
- STDEV
- STDEVP
- STRING_AGG
- SUM
- VAR
- VARP

Reference:
https://docs.microsoft.com/en-us/sql/t-sql/functions/aggregate-functions-transact-sql?view=sql-server-ver1 5

Question: 49

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:

Answer 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.	○	○

Answer:

Answer 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.	○	○

Explanation:
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: 50

HOTSPOT -
Select the answer that correctly completes the sentence.
Hot Area:

Answer 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

Answer:

Answer 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

Explanation:
The same set of columns.

In relational databases, a row is a data record within a table. Each row, which represents a complete record of specific item data, holds different data within the same structure. A row is occasionally referred to as a tuple.

Incorrect:
Not: a key value pair.

Unlike relational databases, key-value databases do not have a specified structure. Relational databases store data in tables where each column has an assigned data type. Key-value databases are a collection of key-value pairs that are stored as individual records and do not have a predefined data structure.

Reference:
https://www.techopedia.com/definition/4425/database-row
https://www.techtarget.com/searchdatamanagement/tip/NoSQL-database-types-explained-Key-value-store

Question: 51

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:

Answer Area

Statements	Yes	No
Normalizing a database increases the throughput of writing transactions.	○	○
Analytics systems are more normalized than transactional systems.	○	○
Normalizing a database results in queries that require more joins.	○	○

Answer:

Answer Area

Statements	Yes	No
Normalizing a database increases the throughput of writing transactions.	○	●
Analytics systems are more normalized than transactional systems.	○	●
Normalizing a database results in queries that require more joins.	●	○

Explanation:
Box 1: No -
Database normalization is the process of restructuring a relational database in accordance with a series of so-called normal forms in order to reduce data redundancy and improve data integrity.

Full normalisation will generally not improve performance, in fact it can often make it worse but it will keep your data duplicate free.

Box 2: No -
Analytics systems are deformalized to increase performance.

Transactional database systems are normalized to increase data consistency.

Box 3: Yes -
Transactional database systems are more normalized and requires more joins.

Reference:
https://www.sqlshack.com/what-is-database-normalization-in-sql-server

Question: 52

HOTSPOT -
Select the answer that correctly completes the sentence.
Hot Area:

Answer Area

▼ natively support the analysis of relationships between entities

- Column family databases
- Document databases
- Graph databases
- Key-value stores

Answer:

Explanation:
A graph database stores two types of information, nodes and edges. Edges specify relationships between nodes. Nodes and edges can have properties that provide information about that node or edge, similar to columns in a table. Edges can also have a direction indicating the nature of the relationship.

Reference:
https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview

Question: 53

HOTSPOT -
Select the answer that correctly completes the sentence.
Hot Area:

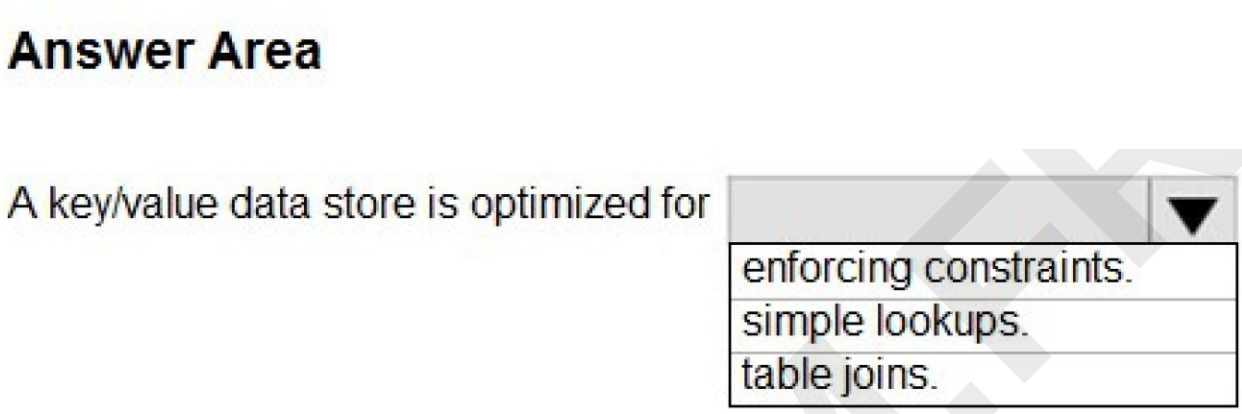

Answer:

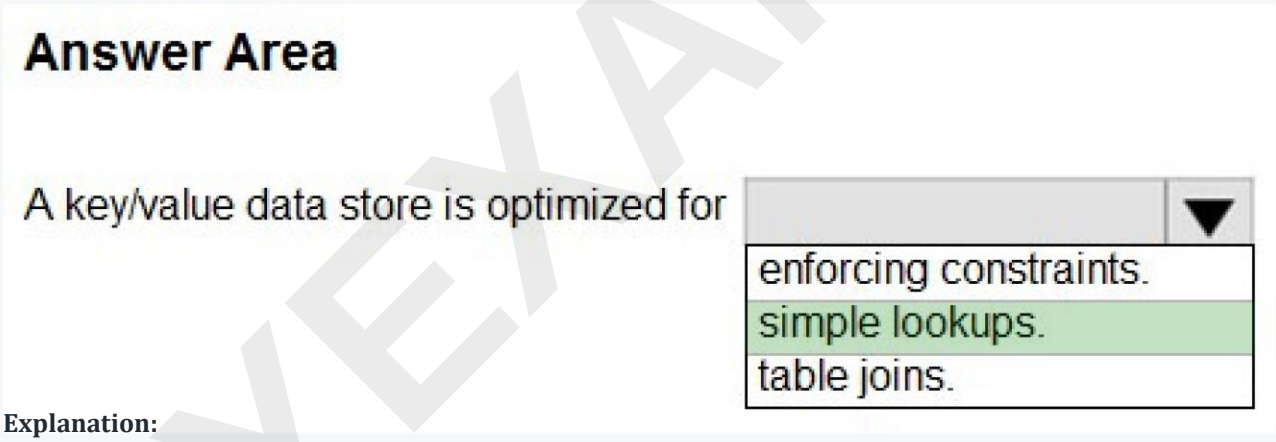

Explanation:
A key/value store associates each data value with a unique key.

Key/value stores are highly optimized for applications performing simple lookups, but are less suitable if you need to query data across different key/value stores.

Key/value stores are also not optimized for querying by value.

A single key/value store can be extremely scalable, as the data store can easily distribute data across multiple nodes on separate machines.

Reference:
https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview

Question: 54

HOTSPOT -
Select the answer that correctly completes the sentence.
Hot Area:

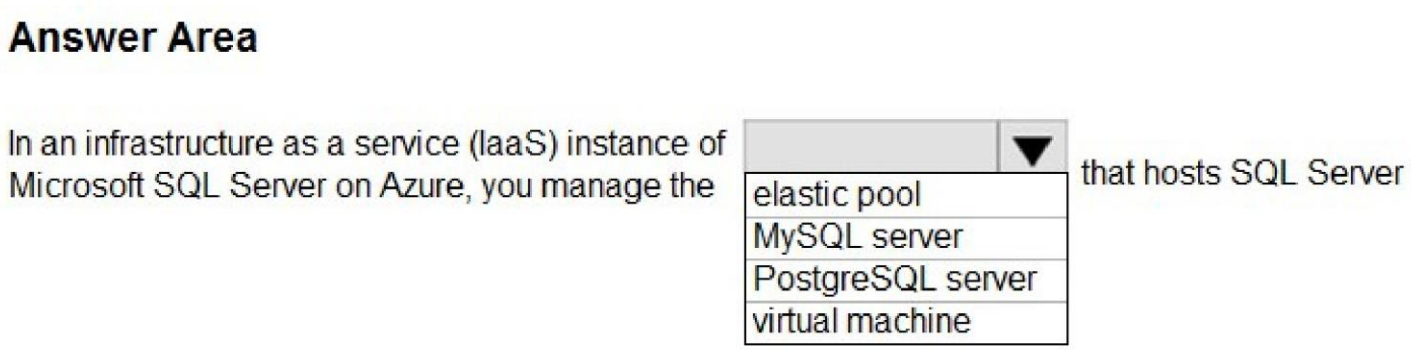

Answer:

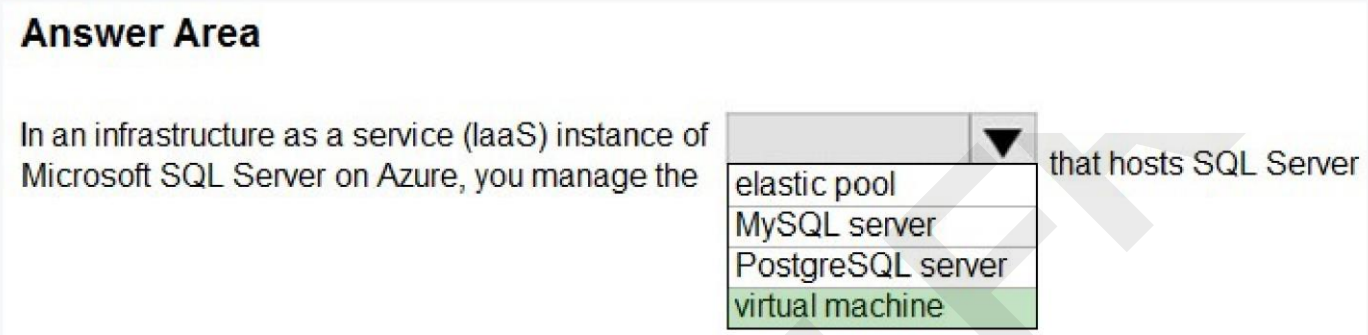

Explanation:
SQL Server on Azure VM -
SQL Server on Azure VM falls into the industry category Infrastructure-as-a-Service (IaaS) and allows you to run SQL Server inside a fully managed virtual machine (VM) in Azure.

The most significant difference from SQL Database and SQL Managed Instance is that SQL Server on Azure Virtual Machines allows full control over the database engine. You can choose when to start maintenance/patching, change the recovery model to simple or bulk-logged, pause or start the service when needed, and you can fully customize the SQL Server database engine. With this additional control comes the added responsibility to manage the virtual machine.

Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/azure-sql-iaas-vs-paas-what-is-overview?view=azuresql

Question: 55

HOTSPOT -
Select the answer that correctly completes the sentence.
Hot Area:

Answer Area

Varying fields for each entity in a JSON document is an example of

- relational data.
- semi-structured data.
- structured data.
- unstructured data.

Answer:

Answer Area

Varying fields for each entity in a JSON document is an example of

- relational data.
- **semi-structured data.**
- structured data.
- unstructured data.

Explanation:
For some semi-structured data formats (e.g. JSON), data sets are frequently a simple concatenation of multiple documents.

Reference:
https://docs.snowflake.com/en/user-guide/semistructured-considerations.html

Question: 56

HOTSPOT -
Select the answer that correctly completes the sentence.
Hot Area:

Answer Area

A JSON document is an example of

- graph data.
- relational data.
- semi-structured data.
- unstructured data.

Answer:

Answer Area

A JSON document is an example of [▼]

- graph data.
- relational data.
- semi-structured data.
- unstructured data.

Explanation:
Semi-structured data -
Semi-structured data (e.g., JSON, CSV, XML) is the bridge between structured and unstructured data. It does not have a predefined data model and is more complex than structured data, yet easier to store than unstructured data.

Reference:
https://www.ibm.com/cloud/blog/structured-vs-unstructured-data

Question: 57

Which property of a transactional workload guarantees that each transaction is treated as a single unit that either succeeds completely or fails completely?

A. atomicity
B. isolation
C. durability
D. consistency

Answer: A

Explanation:
The purpose of making transactions atomic is to prevent different transactions from interfering with one another. This can only happen if more than one user process is trying to access the database at the same time, as when a server allows several clients to use it concurrently. The simplest way to enforce atomicity is for the DBMS to refuse to start any transaction until the previous one has committed.

Reference:
https://www.sciencedirect.com/topics/computer-science/atomic-transaction

Question: 58

You have a data store that has the structure shown in the following exhibit.

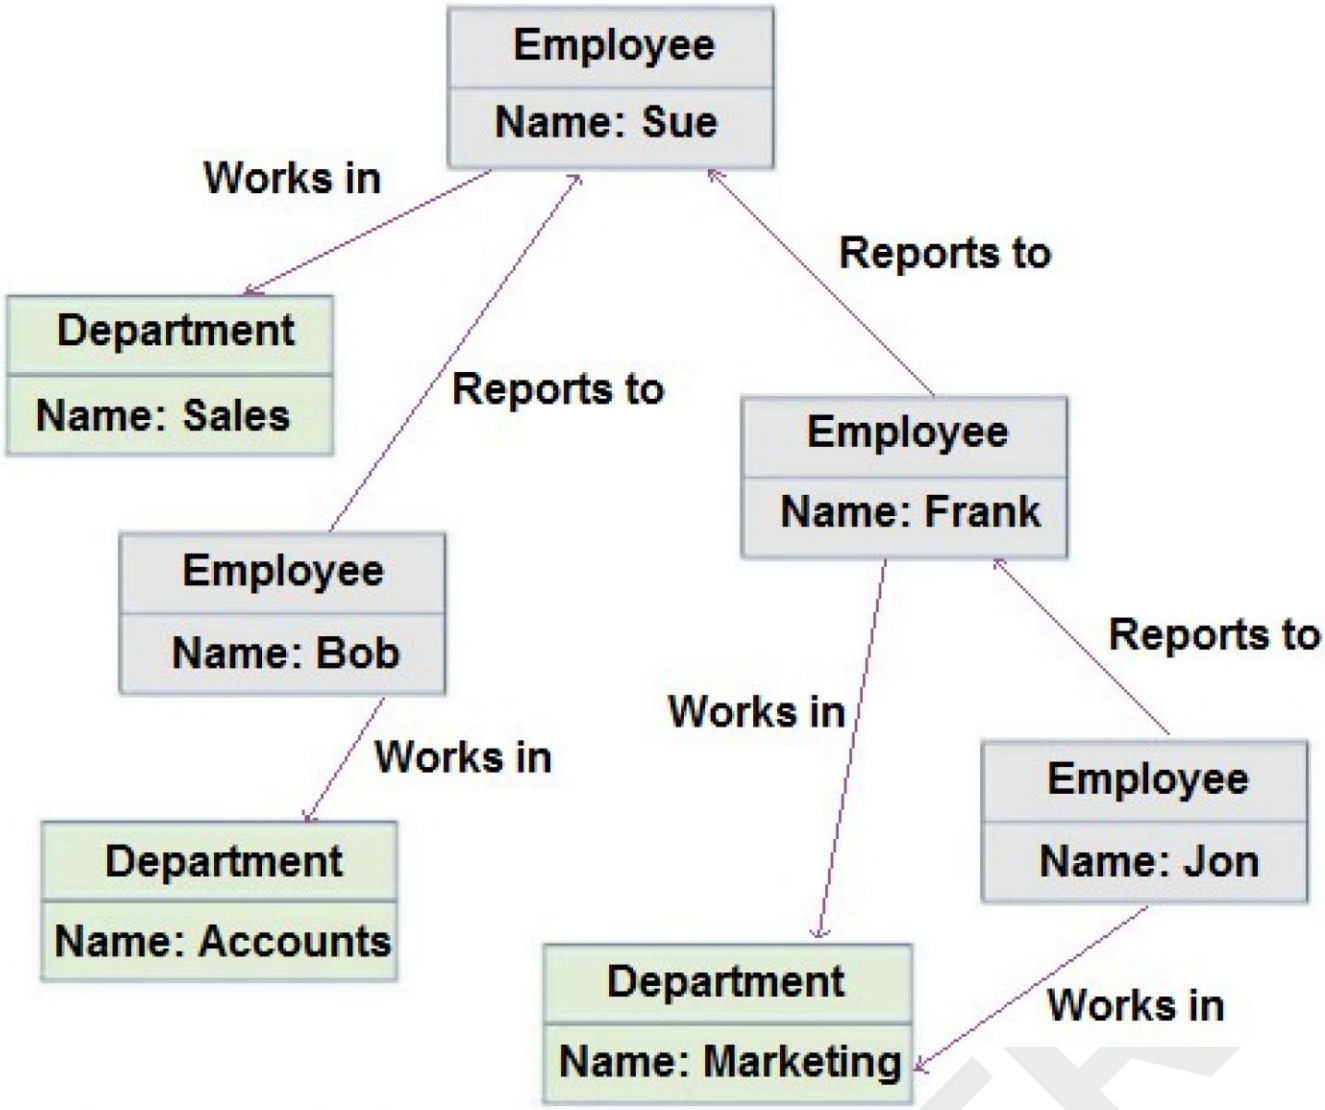

Which type of data store is this?

A. key/value
B. object data
C. graph
D. time series

Answer: C

Explanation:
A graph database stores two types of information, nodes and edges. Edges specify relationships between nodes. Nodes and edges can have properties that provide information about that node or edge, similar to columns in a table. Edges can also have a direction indicating the nature of the relationship.

Reference:
https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview

Question: 59

Which type of database contains nodes and edges?

A. graph
B. key/value

C. columnar

D. time series

Answer: A

Explanation:
A graph database stores two types of information, nodes and edges. Edges specify relationships between nodes. Nodes and edges can have properties that provide information about that node or edge, similar to columns in a table. Edges can also have a direction indicating the nature of the relationship.

Reference:
https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview

Question: 60

HOTSPOT -
Select the answer that correctly completes the sentence.
Hot Area:

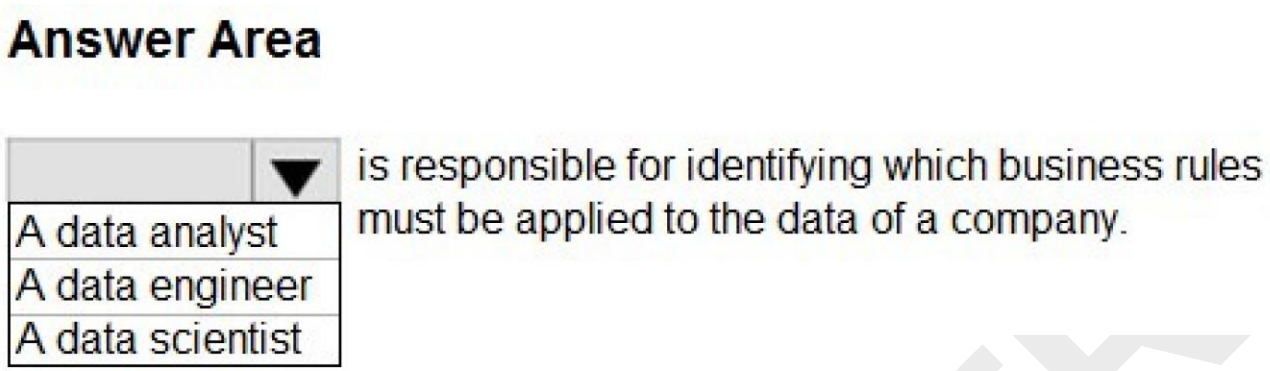

Answer:

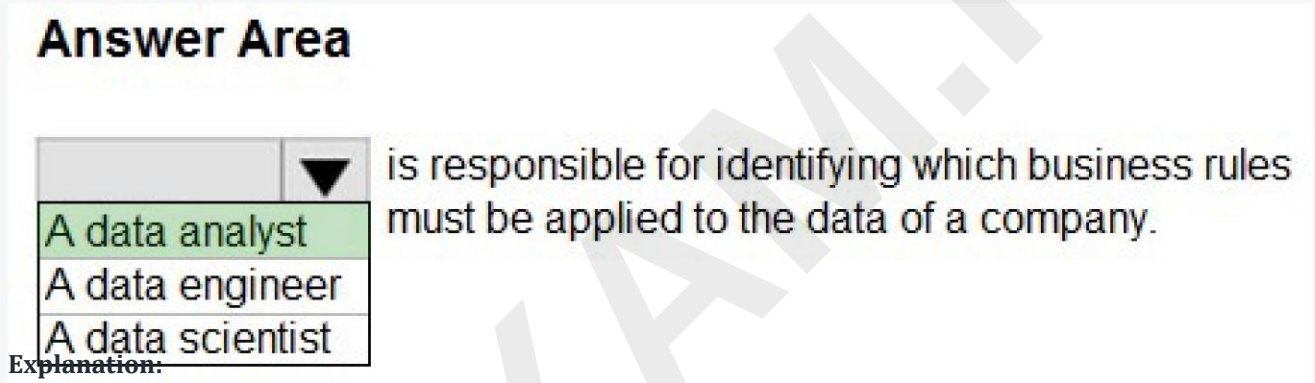

Explanation:
A data analyst's primary skill set revolves around data acquisition, handling, and processing.

Incorrect:
* A data engineer requires an intermediate level understanding of programming to build thorough algorithms along with a mastery of statistics and math.

* A data scientist needs to be a master of both worlds. Data, stats, and math along with in-depth programming knowledge for Machine Learning and Deep
Learning.

Reference:
https://www.edureka.co/blog/data-analyst-vs-data-engineer-vs-data-scientist/

Question: 61

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

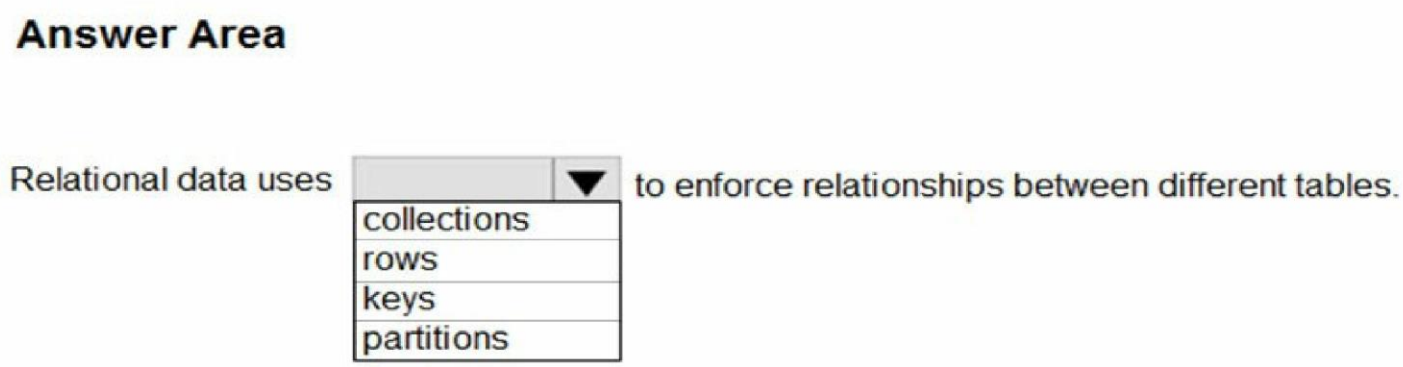

Answer:

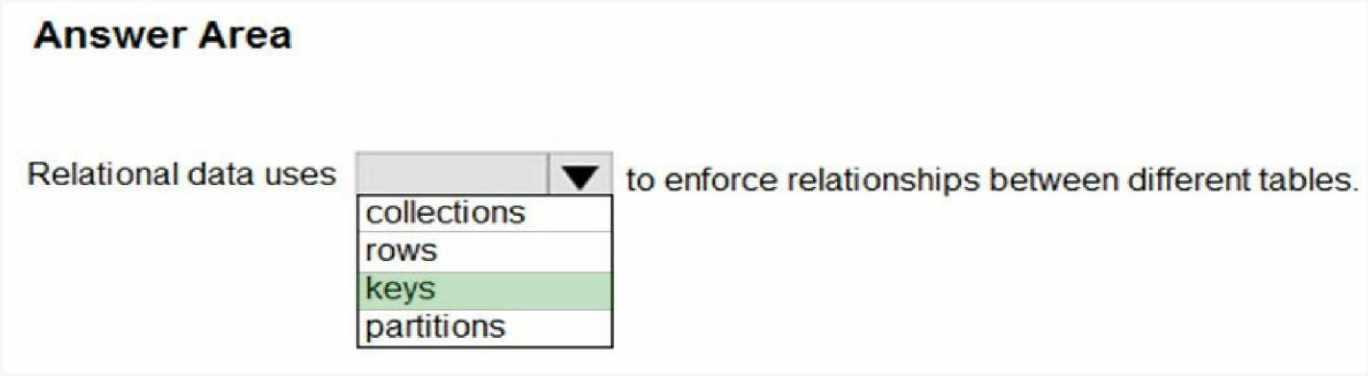

Explanation:

Reference:
https://teachcomputerscience.com/relational-databases/

Question: 62

You have an inventory management database that contains the following table.

ProductName	Quantity
Product1	100
Product2	129
Product3	176

Which statement should you use in a SQL query to change the inventory quantity of Product1 to 270?

A. INSERT
B. MERGE
C. UPDATE
D. CREATE

Answer: C

Explanation:

ALTER - add/delete/modify column

UPDATE - update existing row

INSERT - Insert new row

Reference:

https://docs.microsoft.com/en-us/sql/t-sql/queries/update-transact-sql?view=sql-server-ver15

Question: 63 Your company needs to implement a relational database in Azure. The solution must minimize ongoing maintenance.
Which Azure service should you use?

A. Azure HDInsight
B. Azure SQL Database
C. Azure Cosmos DB
D. SQL Server on Azure Virtual Machines

Answer: B
Explanation:

Reference:
https://azure.microsoft.com/en-us/services/sql-database/#features

Question: 64 You are writing a set of SQL queries that administrators will use to troubleshoot an Azure SQL database. You need to embed documents and query results into a SQL notebook.
What should you use?

A. Microsoft SQL Server Management Studio (SSMS)
B. Azure Data Studio
C. Azure CLI
D. Azure PowerShell

Answer: B
Explanation:
B. Azure Data Studio is a cross-platform database tool that allows you to embed documents and query results into a SQL notebook. It supports Azure SQL databases, SQL Server, and other databases. It also has built-in support for SQL Notebooks, which allows you to mix markdown, code, and results in a single notebook. This makes it a great tool for creating documentation and troubleshooting guides for administrators to use when working with Azure SQL databases.

Reference:
https://www.mssqltips.com/sqlservertip/5997/create-sql-server-notebooks-in-azure-data-studio/

Question: 65

DRAG DROP -
Match the terms to the appropriate descriptions.
To answer, drag the appropriate term from the column on the left to its description on the right. Each term may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.
Select and Place:

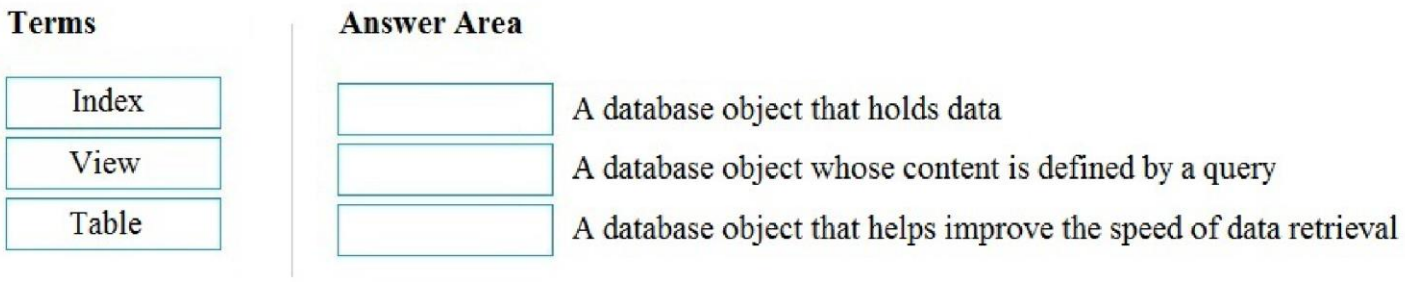

Answer:

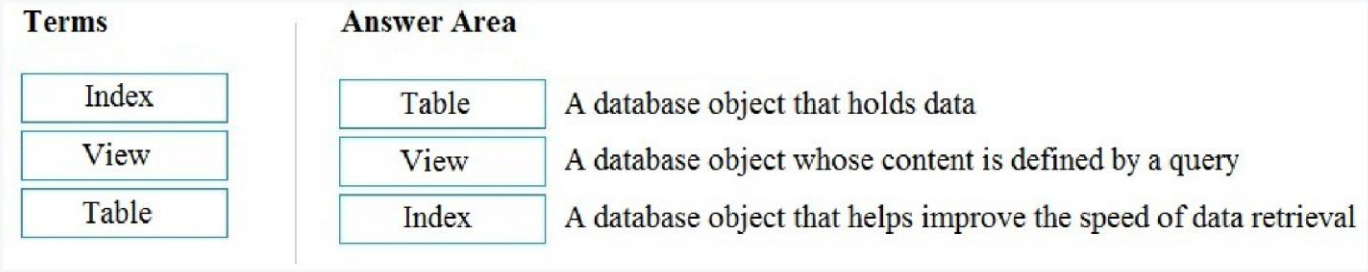

Explanation:

Reference:
https://en.wikipedia.org/wiki/Table_(database)
https://en.wikipedia.org/wiki/View_(SQL)
https://en.wikipedia.org/wiki/Database_index

Question: 66

You have an e-commerce application that reads and writes data to an Azure SQL database. Which type of processing does the application use?

A. stream processing
B. batch processing
C. Online Analytical Processing (OLAP)
D. Online Transaction Processing (OLTP)

Answer: D

Explanation:

OLTP is designed to serve as a persistent data store for business or front-end applications. OLTP administers day to day transaction of an organization.

Online Transaction Processing (OLTP)

https://learn.microsoft.com/en-us/training/modules/explore-core-data-concepts/5-transactional-data-processing?ns-enrollment-type=learningpath&ns-enrollment-id=learn.wwl.azure-data-fundamentals-explore-

core-data-concepts

Reference:

https://sqlwizard.blog/2020/03/15/sql-server-oltp-vs-olap/

Question: 67

When can you use an Azure Resource Manager template?

A. to automate the creation of an interdependent group of Azure resources in a repeatable way
B. to apply Azure policies for multi-tenant deployments
C. to provision Azure subscriptions
D. to control which services and feature administrators and developers can deploy from the Azure portal

Answer: A

Explanation:
You can automate deployments and use the practice of infrastructure as code. In code, you define the infrastructure that needs to be deployed
To implement infrastructure as code for your Azure solutions, use Azure Resource Manager templates (ARM templates). The template is a JavaScript Object
Notation (JSON) file that defines the infrastructure and configuration for your project. The template uses declarative syntax, which lets you state what you intend to deploy without having to write the sequence of programming commands to create it. In the template, you specify the resources to deploy and the properties for those resources.

Reference:
https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/overview

Question: 68

You have an Azure SQL database that you access directly from the Internet.
You recently changed the public IP address of your computer.
After changing the IP address, you can no longer access the database. You can connect to other resources in Azure.
What is a possible cause of the issue?

A. role-based access control (RBAC)
B. Dynamic Host Configuration Protocol (DHCP)
C. Domain Name Service (DNS)
D. a database-level firewall

Answer: D

Explanation:
The Azure SQL Database firewall lets you decide which IP addresses may or may not have access to either your Azure SQL Server or your Azure SQL database.

When creating an Azure SQL Database, the firewall needs to be configured before anyone will be able to access the database. By default, no external access to your SQL Database will be allowed until you explicitly assign permission by creating a firewall rule.

Reference:
https://www.sqlshack.com/configuring-the-azure-sql-database-firewall/

Question: 69

DRAG DROP -
Match the tools to the appropriate descriptions.
To answer, drag the appropriate tool from the column on the left to its description on the right. Each tool may be used once, more than once, or not at all.
Select and Place:

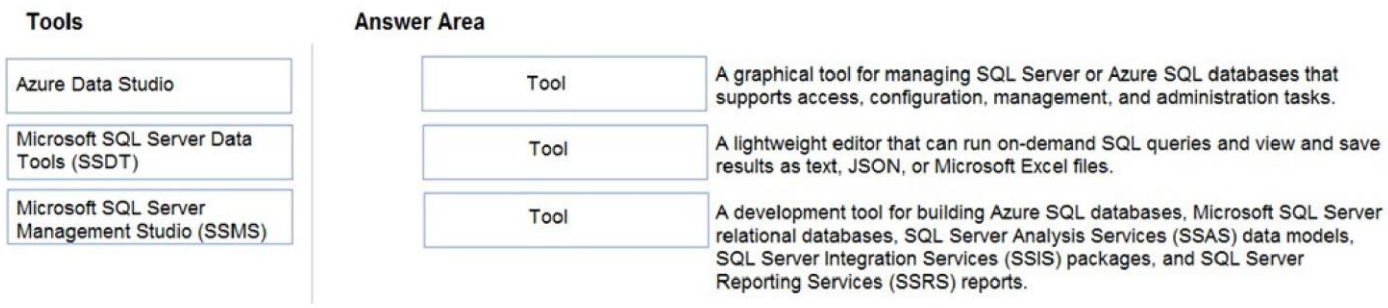

Answer:

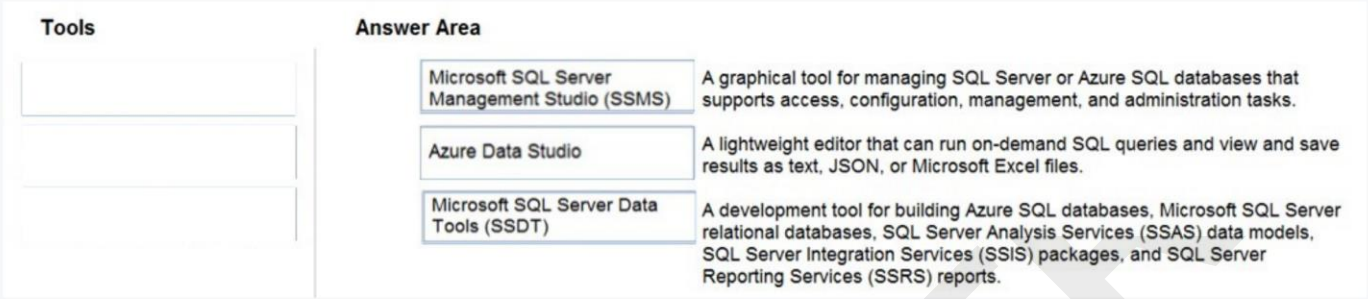

Explanation:
Box 1: Microsoft SQL Server Management Studio (SSMS)
SQL Server Management Studio (SSMS) is an integrated environment for managing any SQL infrastructure, from SQL Server to Azure SQL Database.

Box 2: Azure Data Studio -
Azure Data Studio offers a modern, keyboard-focused SQL coding experience that makes your everyday tasks easier with built-in features, such as multiple tab windows, a rich SQL editor, IntelliSense, keyword completion, code snippets, code navigation, and source control integration (Git). Run on-demand SQL queries, view and save results as text, JSON, or Excel. Edit data, organize your favorite database connections, and browse database objects in a familiar object browsing experience.

Box 3: Microsoft SQL Server Data Tools (SSDT)
SQL Server Data Tools (SSDT) is a modern development tool for building SQL Server relational databases, databases in Azure SQL, Analysis Services (AS) data models, Integration Services (IS) packages, and Reporting Services (RS) reports. With SSDT, you can design and deploy any SQL Server content type with the same ease as you would develop an application in Visual Studio.

Reference:
https://docs.microsoft.com/en-us/sql/ssms/download-sql-server-management-studio-ssms https://code.visualstudio.com/docs/supporting/FAQ https://docs.microsoft.com/en-us/sql/azure-data-studio/what-is-azure-data-studio https://docs.microsoft.com/en-us/sql/ssdt/download-sql-server-data-tools-ssdt

Question: 70

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:

Answer Area

Statements	Yes	No
Relational database tables contain columns and rows	○	○
Indexes in a relational database describe the data types in a table	○	○
A database view is a virtual table whose content is defined by a query	○	○

Answer:

Answer Area

Statements	Yes	No
Relational database tables contain columns and rows	○	○
Indexes in a relational database describe the data types in a table	○	○
A database view is a virtual table whose content is defined by a query	○	○

Explanation:
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: 71

Which command-line tool can you use to query Azure SQL databases?

A. sqlcmd
B. bcp
C. azdata
D. Azure CLI

Answer: A

Explanation:
The sqlcmd utility lets you enter Transact-SQL statements, system procedures, and script files at the command prompt.

Incorrect Answers:
B: The bulk copy program utility (bcp) bulk copies data between an instance of Microsoft SQL Server and a data file in a user-specified format.

D: The Azure CLI is the defacto tool for cross-platform and command-line tools for building and managing Azure resources.

Reference:
https://docs.microsoft.com/en-us/sql/tools/overview-sql-tools?view=sql-server-ver15

Question: 72

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:

Answer Area

Statements	Yes	No
Azure SQL Database includes a managed backup service.	○	○
Azure SQL Database has built-in high availability.	○	○
Azure SQL Database can use Azure Defender.	○	○

Answer:

Answer Area

Statements	Yes	No
Azure SQL Database includes a managed backup service.	○	○
Azure SQL Database has built-in high availability.	○	○
Azure SQL Database can use Azure Defender.	○	○

Explanation:
Box 1: Yes -

Box 2: Yes -

Box 3: Yes -
Azure Defender provides security alerts and advanced threat protection for virtual machines, SQL databases, containers, web applications, your network, and more.

Azure Defender provides security alerts and advanced threat protection for virtual machines, SQL databases, containers, web applications, your network, and more.

Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview https://azure.microsoft.com/en-us/blog/announcing-sql-atp-and-sql-vulnerability-assessment-general-availability/ https://docs.microsoft.com/en-us/azure/security-center/azure-defender

Question: 73

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:

Answer Area

Statements	Yes	No
You can use Azure Data Studio to query a Microsoft SQL Server big data cluster.	○	○
You can use Microsoft SQL Server Management Studio (SSMS) to query an Azure Synapse Analytics data warehouse.	○	○
You can use MySQL Workbench to query Azure Database for MariaDB databases.	○	○

Answer:

Answer Area

Statements	Yes	No
You can use Azure Data Studio to query a Microsoft SQL Server big data cluster.	○	○
You can use Microsoft SQL Server Management Studio (SSMS) to query an Azure Synapse Analytics data warehouse.	○	○
You can use MySQL Workbench to query Azure Database for MariaDB databases.	○	○

Explanation:

A) **Yes** from https://docs.microsoft.com/en-us/sql/azure-data-studio/what-is?view=sql-server-ver15 Use Azure Data Studio if you: Are connecting to a SQL Server 2019 big data cluster

B) **Yes** from https://docs.microsoft.com/en-us/learn/modules/query-azure-sql-data-warehouse/4-query-dw-using-ssms

C) **Yes** from https://docs.microsoft.com/en-us/azure/mariadb/connect-workbench

Reference:

https://docs.microsoft.com/en-us/sql/big-data-cluster/connect-to-big-data-cluster?view=sql-server-ver15

https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-query-ssms

https://docs.microsoft.com/en-us/azure/mariadb/connect-workbench

Question: 74 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:

Answer Area

Statements	Yes	No
Platform as a service (PaaS) database offerings in Azure provide built-in high availability.	○	○
Platform as a service (PaaS) database offerings in Azure provide configurable scaling options.	○	○
Platform as a service (PaaS) database offerings in Azure reduce the administrative overhead for managing hardware.	○	○

Answer:

Answer Area

Statements	Yes	No
Platform as a service (PaaS) database offerings in Azure provide built-in high availability.	○	○
Platform as a service (PaaS) database offerings in Azure provide configurable scaling options.	○	○
Platform as a service (PaaS) database offerings in Azure reduce the administrative overhead for managing hardware.	○	○

Explanation:

Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview

Question: 75

HOTSPOT -
You have the following SQL query.

```
INSERT INTO dbo.Products (ProductID, ProductName, Price, ProductDescription)
    VALUES (1, 'Clamp', 12.48, 'Workbench clamp') ;
```

What are dbo.Products and ProductName? To answer, select the appropriate options in the answer area. NOTE:
Each correct selection is worth one point.
Hot Area:

Answer Area

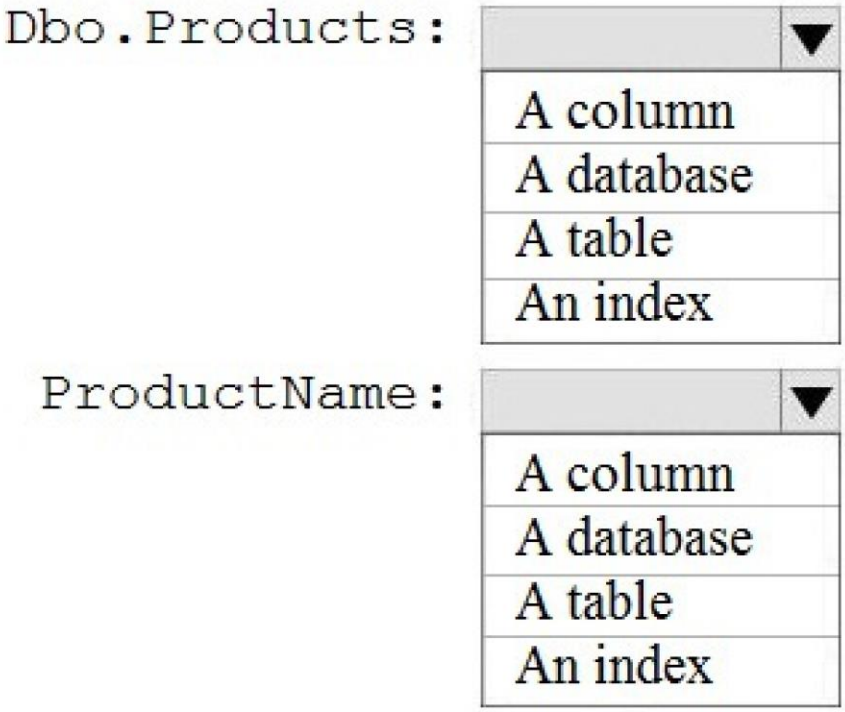

Answer:

Answer Area

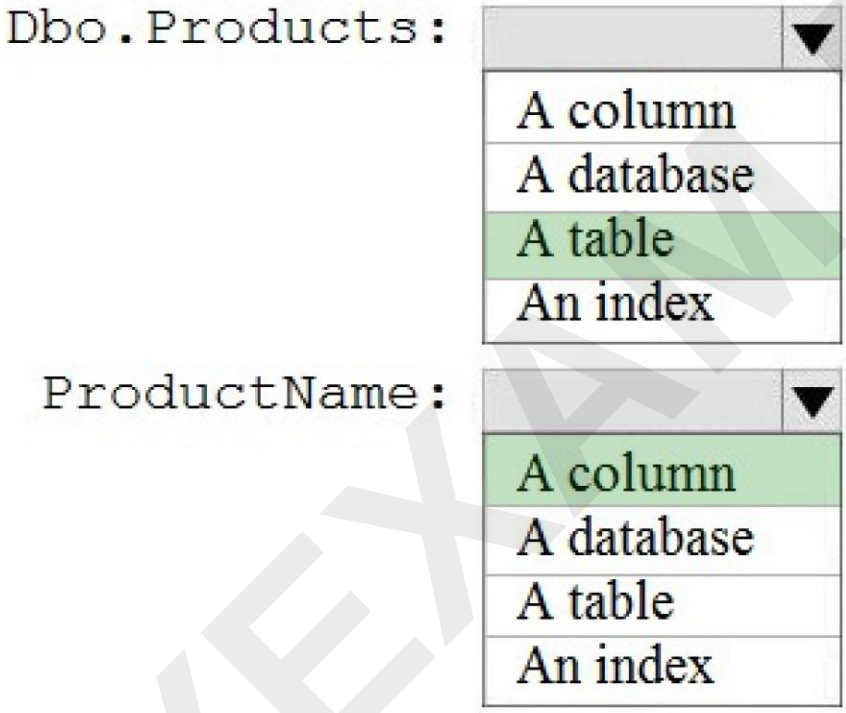

Question: 76

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:

Answer Area

Statements	Yes	No
You must apply operating system updates to Azure SQL databases regularly.	○	○
You need a Microsoft 365 subscription to create an Azure SQL database.	○	○
You can use existing Microsoft SQL Server licenses to reduce the cost of Azure SQL databases.	○	○

Answer:

Answer Area

Statements	Yes	No
You must apply operating system updates to Azure SQL databases regularly.	○	○
You need a Microsoft 365 subscription to create an Azure SQL database.	○	○
You can use existing Microsoft SQL Server licenses to reduce the cost of Azure SQL databases.	○	○

Explanation:

N/N/Y

Can I bring my on-premises or IaaS SQL Server license to Hyperscale? Yes, Azure Hybrid Benefit ...

(https://docs.microsoft.com/en-us/azure/azure-sql/database/service-tier-hyperscale-frequently-asked-questions-faq

https://azure.microsoft.com/en-us/pricing/hybrid-benefit/#calculator)

Azure Hybrid Benefit allows you to use SQL Server licenses with Software Assurance or qualifying subscription licenses to pay a reduced base rate* for these products and services for SQL Server on Azure: vCPU-based service tiers of Azure SQL Database (excluding serverless).

Azure SQL Managed Instance.

SQL Server in Azure Virtual Machines.

SQL Server Integration Services.

Reference:

https://azure.microsoft.com/en-gb/blog/hot-patching-sql-server-engine-in-azure-sql-database/
https://azure.microsoft.com/en-us/services/sql-database/#product-overview

Question: 77

Which statement is an example of Data Definition Language (DDL)?

A. SELECT
B. JOIN
C. UPDATE
D. CREATE

Answer: D

Explanation:
Data Definition Language (DDL) statements defines data structures. Use these statements to create, alter, or drop data structures in a database. These statements include:
☞ ALTER
☞ Collations
☞ CREATE
☞ DROP
☞ DISABLE TRIGGER
☞ ENABLE TRIGGER
☞ RENAME
☞ UPDATE STATISTICS
☞ TRUNCATE TABLE
☞ INSERT

Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/statements

Question: 78

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:

Answer Area

Statements	Yes	No
Azure Data Studio can be used to query an Azure SQL database from a device that runs macOS.	○	○
Microsoft SQL Server Management Studio (SSMS) enables users to create and use SQL notebooks.	○	○
Azure Data Studio can be used to restore a database.	○	○

Answer:

Answer Area

Statements	Yes	No
Azure Data Studio can be used to query an Azure SQL database from a device that runs macOS.	○	○
Microsoft SQL Server Management Studio (SSMS) enables users to create and use SQL notebooks.	○	○
Azure Data Studio can be used to restore a database.	○	○

Explanation:
Box 1: Yes -
Azure Data Studio is a cross-platform database tool for data professionals using on-premises and cloud data platforms on Windows, macOS, and Linux.

You can use Azure Data Studio to connect to an Azure SQL Database server. You'll then run Transact-SQL (T-SQL) statements to create and query Azure SQL databases.

Box 2: No -
SQL Server Management Studio is for configuring, managing, and administering all components within Microsoft SQL Server, not to create SQL notebooks.

Instead use Azure Data Studio to create SQL notebook.

Box 3: Yes -
You can use the Azure Data Studio to restore databases.

Reference:
https://docs.microsoft.com/en-us/sql/azure-data-studio/what-is-azure-data-studio

Question: 79

You are deploying a software as a service (SaaS) application that requires a relational database for Online Transaction Processing (OLTP).
Which Azure service should you use to support the application?

A. Azure Cosmos DB
B. Azure HDInsight
C. Azure SQL Database
D. Azure Synapse Analytics

Answer: C

Explanation:
Azure SQL Database is relational database and a managed service.
Incorrect Answers:
A, B: Cosmos DB, HDInsight are non-relational databases.
D: Azure Synapse Analytics is for data warehousing, not for Online Transaction Processing

Reference:
https://cloud.netapp.com/blog/azure-cvo-blg-azure-database-review-your-guide-for-database-assessment

Question: 80

What are two benefits of platform as a service (PaaS) relational database offerings in Azure, such as Azure SQL Database? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. access to the latest features
B. complete control over backup and restore processes
C. in-database machine learning services
D. reduced administrative effort for managing the server infrastructure

Answer: AD

Explanation:
A: Azure SQL Database is a fully managed platform as a service (PaaS) database engine that handles most of the database management functions such as upgrading, patching, backups, and monitoring without user involvement.
D: SQL Database delivers predictable performance with multiple resource types, service tiers, and compute sizes. It provides dynamic scalability with no downtime, built-in intelligent optimization, global scalability and availability, and advanced security options. These capabilities allow you to focus on rapid app development and accelerating your time-to-market, rather than on managing virtual machines and infrastructure.

Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview