

Microsoft.DP-900.vOct-2024.by.Lino.110q

Number: DP-900
Passing Score: 800
Time Limit: 120
File Version: 12.0

Exam Code: DP-900
Exam Name: Microsoft Azure Data Fundamentals



Exam A

QUESTION 1

HOTSPOT

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

Hot Area:

Answer Area

You can query a graph database in Azure Cosmos DB

- as a JSON document by using a SQL-like language.
- as a partitioned row store by using Cassandra Query Language (CQL).
- as a partitioned row store by using Language-Integrated Query (LINQ).
- as nodes and edges by using the Gremlin language.

Answer Area:

Answer Area

You can query a graph database in Azure Cosmos DB

- as a JSON document by using a SQL-like language.
- as a partitioned row store by using Cassandra Query Language (CQL).
- as a partitioned row store by using Language-Integrated Query (LINQ).
- as nodes and edges by using the Gremlin language.

Section:

Explanation:

Reference:

<https://www.sqlshack.com/graph-database-implementation-with-azure-cosmos-db-using-the-api/>

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
The Azure Cosmos DB API is configured separately for each database in an Azure Cosmos DB account.	<input type="radio"/>	<input type="radio"/>
Partition keys are used in Azure Cosmos DB to optimize queries.	<input type="radio"/>	<input type="radio"/>
Items contained in the same Azure Cosmos DB logical partition can have different partition keys.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
The Azure Cosmos DB API is configured separately for each database in an Azure Cosmos DB account.	<input type="radio"/>	<input checked="" type="radio"/>
Partition keys are used in Azure Cosmos DB to optimize queries.	<input checked="" type="radio"/>	<input type="radio"/>
Items contained in the same Azure Cosmos DB logical partition can have different partition keys.	<input type="radio"/>	<input checked="" type="radio"/>

Section:

Explanation:

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>

QUESTION 3

HOTSPOT

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

Hot Area:

Answer Area

A block of code that runs in a database is called

<input type="text"/>
a stored procedure.
a table.
a view.
an index.

Answer Area:

Answer Area

A block of code that runs in a database is called

<input type="text"/>
a stored procedure.
a table.
a view.
an index.

Section:

Explanation:

Reference:

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

QUESTION 4

HOTSPOT

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

Hot Area:



Answer Area:



Section:

Explanation:

QUESTION 5

When provisioning an Azure Cosmos DB account, which feature provides redundancy within an Azure region?

- A. multi-master replication
- B. Availability Zones
- C. the strong consistency level
- D. automatic failover

Correct Answer: B

Section:

Explanation:

Section: Describe how to work with non-relational data on Azure

With Availability Zone (AZ) support, Azure Cosmos DB will ensure replicas are placed across multiple zones within a given region to provide high availability and resiliency to zonal failures.

Note: Azure Cosmos DB provides high availability in two primary ways. First, Azure Cosmos DB replicates data across regions configured within a Cosmos account. Second, Azure Cosmos DB maintains 4 replicas of data within a region.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/high-availability>

QUESTION 6

What is a benefit of the Azure Cosmos DB Table API as compared to Azure Table storage?

- A. provides resiliency if an Azure region fails
- B. supports partitioning
- C. provides a higher storage capacity
- D. supports a multi-master model

Correct Answer: D

Section:

Explanation:

Section: Describe how to work with non-relational data on Azure

Multi-master support for Azure Cosmos DB is now available in all public regions.

Azure CosmosDB table API is a key-value storage hosted in the cloud. It's a part of Azure Cosmos DB, that is Microsoft's multi-model database.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/table-support>

QUESTION 7

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

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

Correct Answer: D

Section:

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 8

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

- A. SELECT
- B. INSERT
- C. DELETE
- D. DROP

Correct Answer: D

Section:



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

Reference:

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

QUESTION 9

A team of developers has computers that run Windows 10 and Ubuntu Desktop.

The developers need to connect to and query an Azure SQL database from each of their computers. The developers require code assistance features such as IntelliSense.

What should the developers use?

- A. sqlcmd
- B. Microsoft SQL Server Management Studio (SSMS)
- C. Azure Data Studio
- D. Azure Data Explorer

Correct Answer: C

Section:

Explanation:

Azure Data Studio is a cross-platform database tool for data professionals who use on-premises and cloud data platforms on Windows, macOS, and Linux.

Azure Data Studio offers a modern editor experience with IntelliSense, code snippets, source control integration, and an integrated terminal.

Reference:

<https://docs.microsoft.com/en-us/sql/azure-data-studio/download-azure-data-studio>

QUESTION 10

You need to store data by using Azure Table storage.

What should you create first?

- A. an Azure Cosmos DB instance
- B. a storage account
- C. a blob container
- D. a table

Correct Answer: B

Section:

Explanation:

First create an Azure storage account, then use Table service in the Azure portal to create a table.

Note: An Azure storage account contains all of your Azure Storage data objects: blobs, files, queues, and tables.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/tables/table-storage-quickstart-portal>

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-create>



QUESTION 11

You need to recommend a data store service that meets the following requirements:

Native SQL API access

Configurable indexes

What should you recommend?

- A. Azure Files
- B. Azure Blob storage
- C. Azure Table storage
- D. Azure Cosmos DB

Correct Answer: D

Section:

Explanation:

Azure Cosmos DB comes with native Core (SQL) API support.

In Azure Cosmos DB, data is indexed following indexing policies that are defined for each container. The default indexing policy for newly created containers enforces range indexes for any string or number. This policy can be overridden with your own custom indexing policy.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-manage-indexing-policy>

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. graph
- B. key/value
- C. document
- D. columnar



Correct Answer: A

Section:

Explanation:

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

What are three characteristics of an Online Transaction Processing (OLTP) workload? Each correct answer presents a complete solution. (Choose three.)

NOTE: Each correct selection is worth one point.

- A. denormalized data
- B. heavy writes and moderate reads
- C. light writes and heavy reads
- D. schema defined in a database
- E. schema defined when reading unstructured data from a database

F. normalized data

Correct Answer: B, D, F

Section:

Explanation:

B: Transactional data tends to be heavy writes, moderate reads.

D: Typical traits of transactional data include: schema on write, strongly enforced. The schema is defined in a database.

F: Transactional data tends to be highly normalized.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/online-transaction-processing>

QUESTION 14

What is the primary purpose of a data warehouse?

- A. to provide answers to complex queries that rely on data from multiple sources
- B. to provide transformation services between source and target data stores
- C. to provide read-only storage of relational and non-relational historical data
- D. to provide storage for transactional line-of-business (LOB) applications

Correct Answer: C

Section:

Explanation:

Consider using a data warehouse when you need to keep historical data separate from the source transaction systems for performance reasons. Data warehouses make it easy to access historical data from multiple locations, by providing a centralized location using common formats, keys, and data models.

Query both relational and nonrelational data.

Incorrect Answers:

D: Data warehouses don't need to follow the same terse data structure you may be using in your OLTP databases.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/data-warehousing>

QUESTION 15

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

Correct Answer: D

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/cli/azure/storage/container?view=azure-cli-latest>

QUESTION 16

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

Correct Answer: B

Section:

QUESTION 17

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

Correct Answer: B

Section:

Explanation:

Reference:

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

QUESTION 18

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

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

Correct Answer: D

Section:

QUESTION 19

Which setting can only be configured during the creation of an Azure Cosmos DB account?

- A. geo-redundancy
- B. multi-region writes
- C. production or non-production account type
- D. API

Correct Answer: D

Section:

Explanation:

You must create a separate account for each API.

Reference:



<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-manage-database-account>

QUESTION 20

You need to perform hybrid transactional and analytical processing (HTAP) queries against Azure Cosmos DB data sources by using Azure Synapse Analytics. What should you use?

- A. Synapse pipelines
- B. a Synapse SQL pool
- C. Synapse Link
- D. Synapse Studio

Correct Answer: C

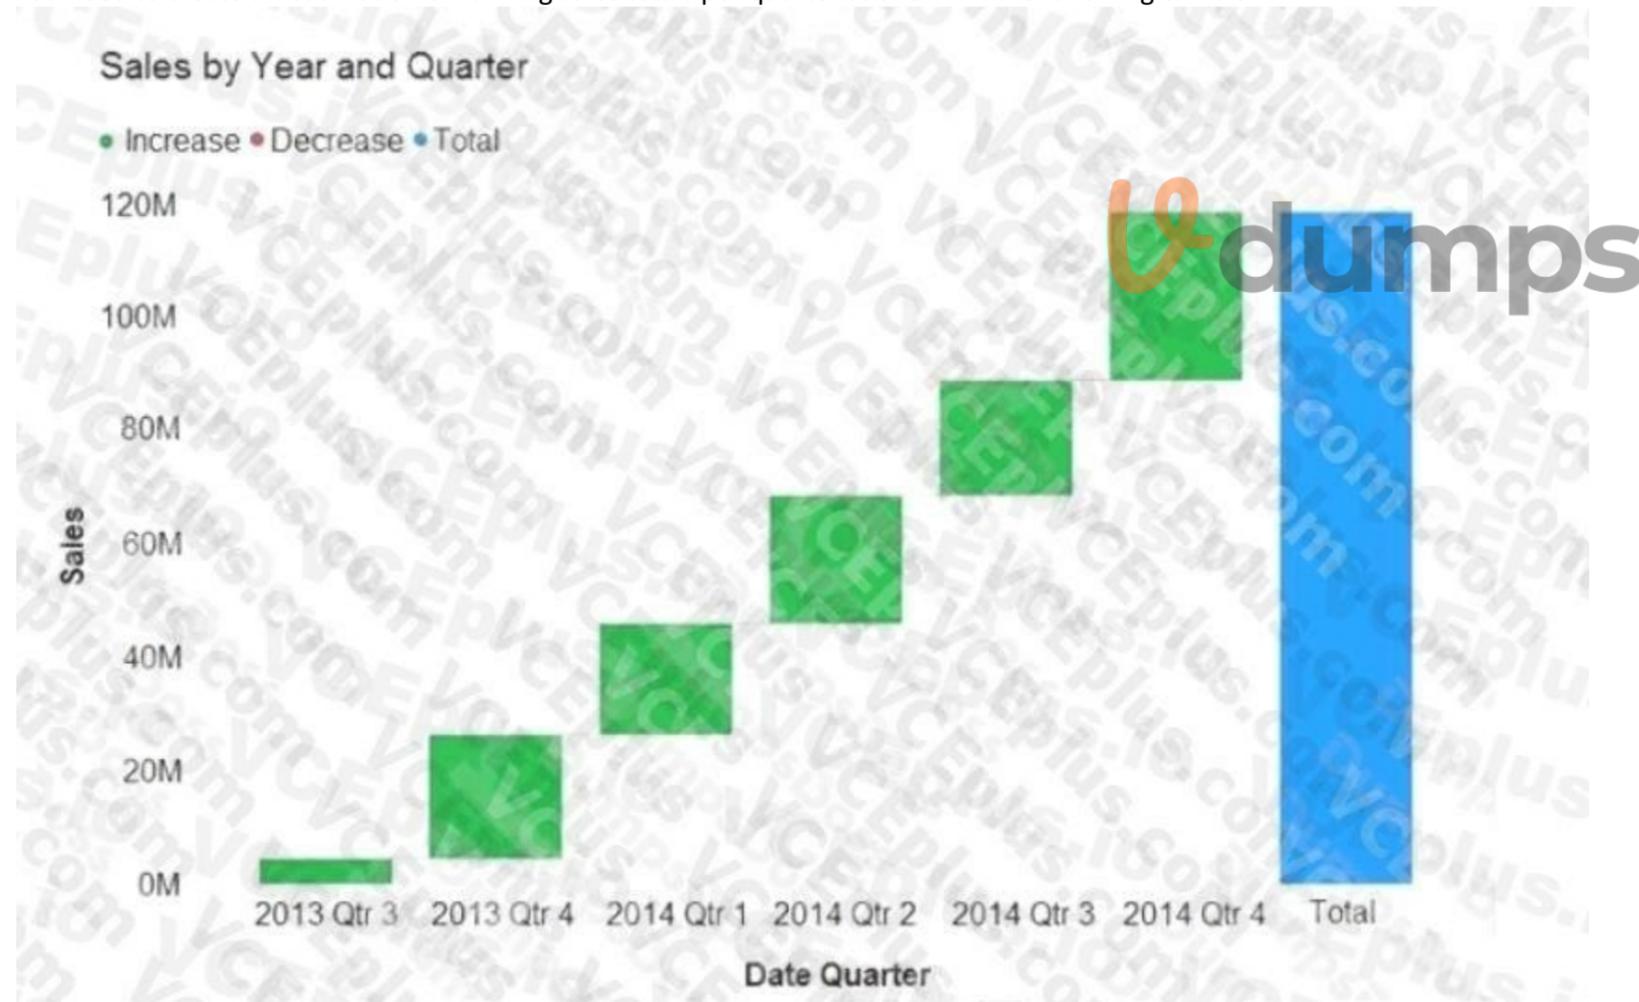
Section:

Explanation:

<https://docs.microsoft.com/en-us/azure/cosmos-db/synapse-link>

QUESTION 21

You need to create a visualization of running sales totals per quarter as shown in the following exhibit.



What should you create in Power BI Desktop?

- A. a waterfall chart
- B. a ribbon chart

- C. a bar chart
- D. a decomposition tree

Correct Answer: A

Section:

Explanation:

<https://docs.microsoft.com/en-us/power-bi/visuals/power-bi-visualization-types-for-reports-and-q-and-a>

QUESTION 22

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

Correct Answer: A

Section:

Explanation:

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

QUESTION 23

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

Section:

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 24

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.

Correct Answer: C

Section:

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 25

SELECT, INSERT, and UPDATE are examples of which type of SQL statement?

- A. Data Definition Language (DDL)
- B. Data Control Language (DCL)
- C. Data Manipulation Language (DML)

Correct Answer: C

Section:

QUESTION 26

Which Azure Storage service implements the key/value model?

- A. Azure Files
- B. Azure Table
- C. Azure Blob
- D. Azure Queue

Correct Answer: B

Section:

QUESTION 27

You have data in a tabular format as shown in the following table.

ID	First name	Last name	Email
1000	Ben	Smith	bensmith@contoso.com
1001	Carlos	Griolo	carlosgriolo@contoso.com

Of which type of data is this an example?

- A. unstructured data
- B. semi-structured data
- C. document
- D. structured data

Correct Answer: D

Section:

QUESTION 28

What can you use to create data ingestion pipelines?



- A. Azure HDInsight
- B. Azure SQL Database
- C. Azure Data Factory
- D. Azure Cosmos DB

Correct Answer: C

Section:

QUESTION 29

Which service is optimized for IoT scenarios that involve streaming time series data?

- A. SQL Server on Azure Virtual Machines
- B. Azure SQL Managed Instance
- C. Azure SQL Database
- D. Azure SQL Edge

Correct Answer: D

Section:

QUESTION 30

What is the responsibility of a database administrator?

- A. applying data cleaning routines
- B. creating data visualizations
- C. managing data integration processes
- D. storing backup copies of data

Correct Answer: C

Section:

QUESTION 31

HOTSPOT

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

Hot Area:

Answer Area

A relational database is appropriate for scenarios that involve a high volume of

- changes to relationships between entities
- geographically distributed writes
- transactional writes
- writes that have varying data structures



Answer Area:

Answer Area

A relational database is appropriate for scenarios that involve a high volume of

- changes to relationships between entities
- geographically distributed writes
- transactional writes
- writes that have varying data structures

Section:

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 32

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	<input type="radio"/>	<input type="radio"/>
Batch processing can output data to a relational database	<input type="radio"/>	<input type="radio"/>
Batch processing can output data to a NoSQL database	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
Batch processing can output data to a file store	<input checked="" type="radio"/>	<input type="radio"/>
Batch processing can output data to a relational database	<input type="radio"/>	<input checked="" type="radio"/>
Batch processing can output data to a NoSQL database	<input type="radio"/>	<input checked="" type="radio"/>

Section:

Explanation:

Box 1: Yes Big data solutions often use long-running batch jobs to filter, aggregate, and otherwise prepare the data for analysis. Usually these jobs involve reading source files from scalable storage (like HDFS, Azure Data Lake Store, and Azure

Storage), processing them, and writing the output to new files in scalable storage.

Box 2: No

Box 3: No

Reference:

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



QUESTION 33

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.

<input type="text"/>	▼
cognitive	
descriptive	
predictive	
prescriptive	

Answer 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

Section:

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 34

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.	<input type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.	<input type="radio"/>	<input type="radio"/>
All relation and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.	<input type="radio"/>	<input type="radio"/>

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.	<input checked="" type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.	<input type="radio"/>	<input checked="" type="radio"/>
All relation and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.	<input type="radio"/>	<input checked="" type="radio"/>

Section:

Explanation:

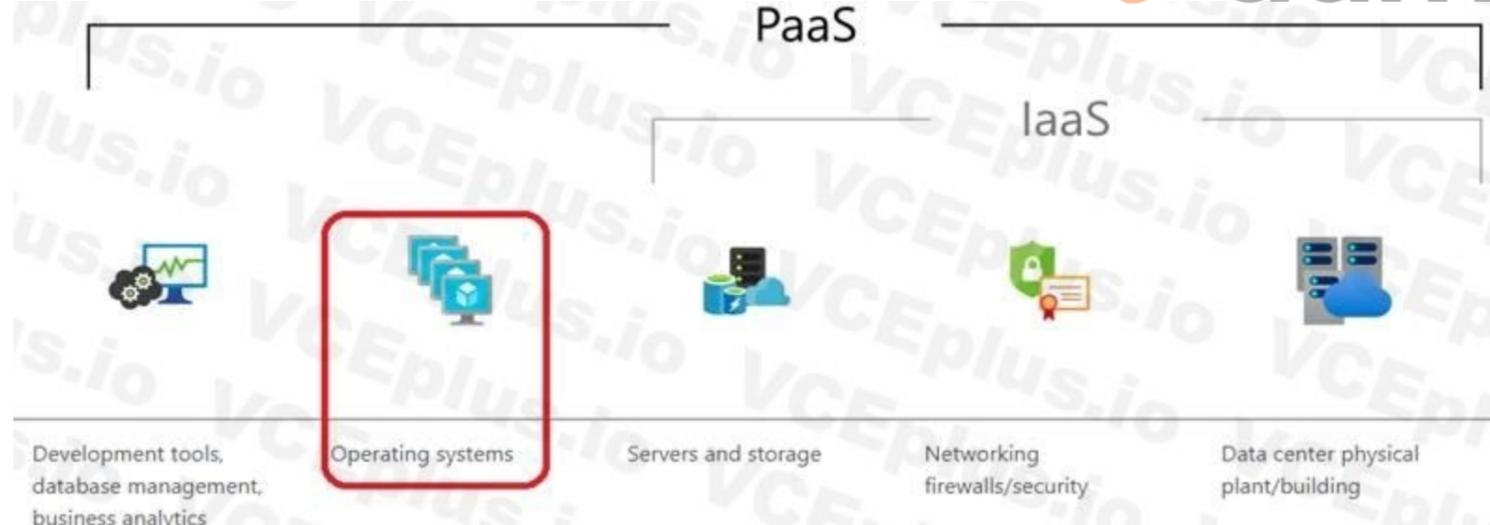
Box 1: Yes

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 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:

Statements	Yes	No
Azure SQL Database includes a fully managed backup service.	<input type="radio"/>	<input type="radio"/>
Azure SQL Database has built-in high availability.	<input type="radio"/>	<input type="radio"/>
Azure SQL Database can use Azure Advanced Threat Protection (ATP).	<input type="radio"/>	<input type="radio"/>

Answer Area:

Statements	Yes	No
Azure SQL Database includes a fully managed backup service.	<input checked="" type="radio"/>	<input type="radio"/>
Azure SQL Database has built-in high availability.	<input checked="" type="radio"/>	<input type="radio"/>
Azure SQL Database can use Azure Advanced Threat Protection (ATP).	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

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/>

QUESTION 36

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.	<input type="radio"/>	<input type="radio"/>
You can use Microsoft SQL Server Management Studio (SSMS) to query an Azure Synapse Analytics data warehouse.	<input type="radio"/>	<input type="radio"/>
You can use MySQL Workbench to query Azure Database for MariaDB databases.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
You can use Azure Data Studio to query a Microsoft SQL Server big data cluster.	<input checked="" type="radio"/>	<input type="radio"/>
You can use Microsoft SQL Server Management Studio (SSMS) to query an Azure Synapse Analytics data warehouse.	<input checked="" type="radio"/>	<input type="radio"/>
You can use MySQL Workbench to query Azure Database for MariaDB databases.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

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 37

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

Azure Data Explorer uses  to explore, filter, and aggregate data quickly.

Answer Area:
Answer Area

Azure Data Explorer uses  to explore, filter, and aggregate data quickly.

Section:
Explanation:

QUESTION 38
DRAG DROP

Match the Azure services to the appropriate requirements.

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

NOTE: Each correct match is worth one point.

Select and Place:

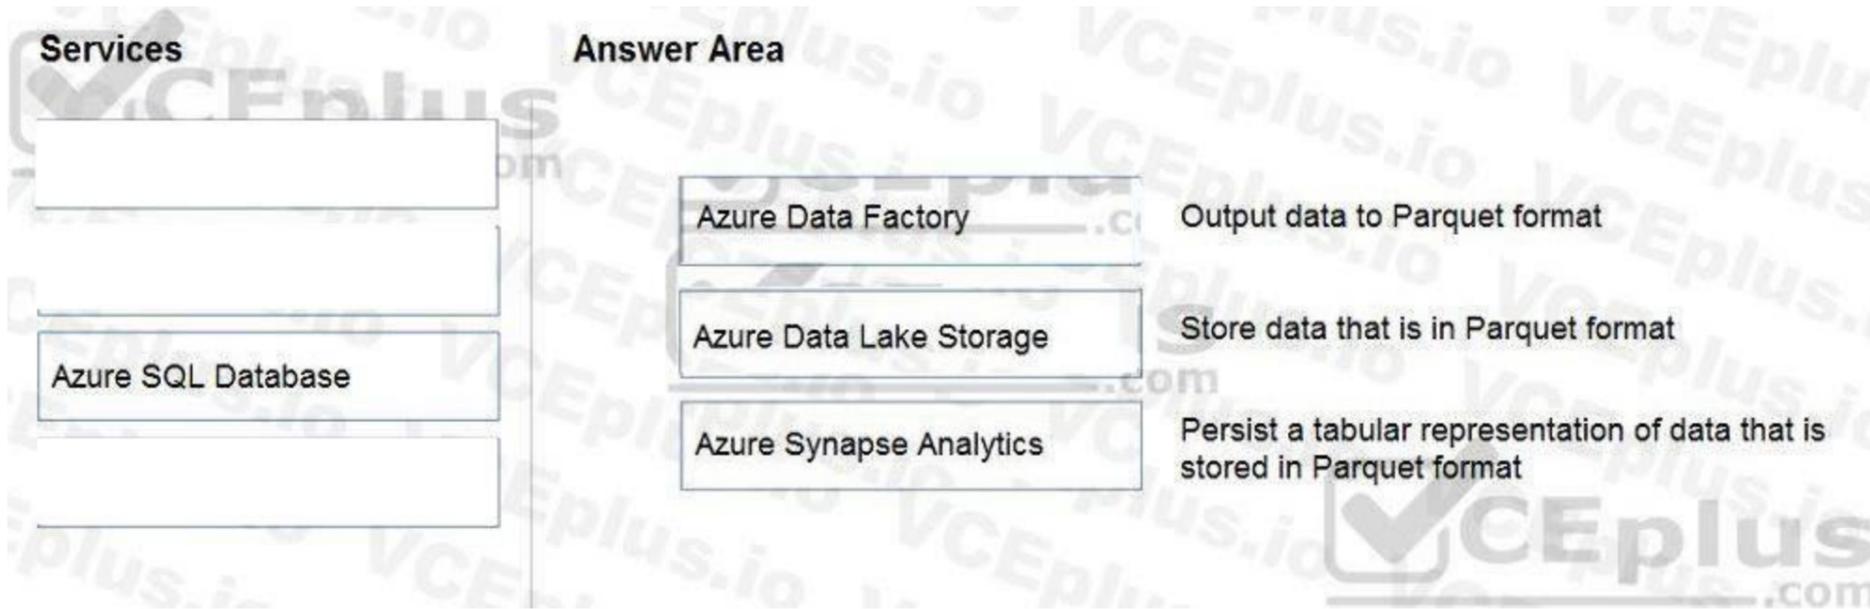
Services

- Azure Data Factory
- Azure Data Lake Storage
- Azure SQL Database
- Azure Synapse Analytics

Answer Area

- Service: Output data to Parquet format
- Service: Store data that is in Parquet format
- Service: Persist a tabular representation of data that is stored in Parquet format

Correct Answer:



Section:

Explanation:

Box 1: Azure Data Factory

Box 2: Azure Data Lake Storage

Azure Data Lake Storage (ADLA) now natively supports Parquet files. ADLA adds a public preview of the native extractor and outputter for the popular Parquet file format

Box 3: Azure Synapse Analytics

Use Azure Synapse Analytics Workspaces.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/supported-file-formats-and-compression-codecs>



QUESTION 39

DRAG DROP

Match the Azure services to the appropriate locations in the architecture.

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

NOTE: Each correct match is worth one point.

Select and Place:

Services

Azure Analysis Services

Azure Data Factory

Azure Table storage

Azure Cosmos DB

Azure Synapse Analytics

Answer Area

SQL Server

SAP HANA

Azure Cosmos DB

Extract Transform Load (ETL)

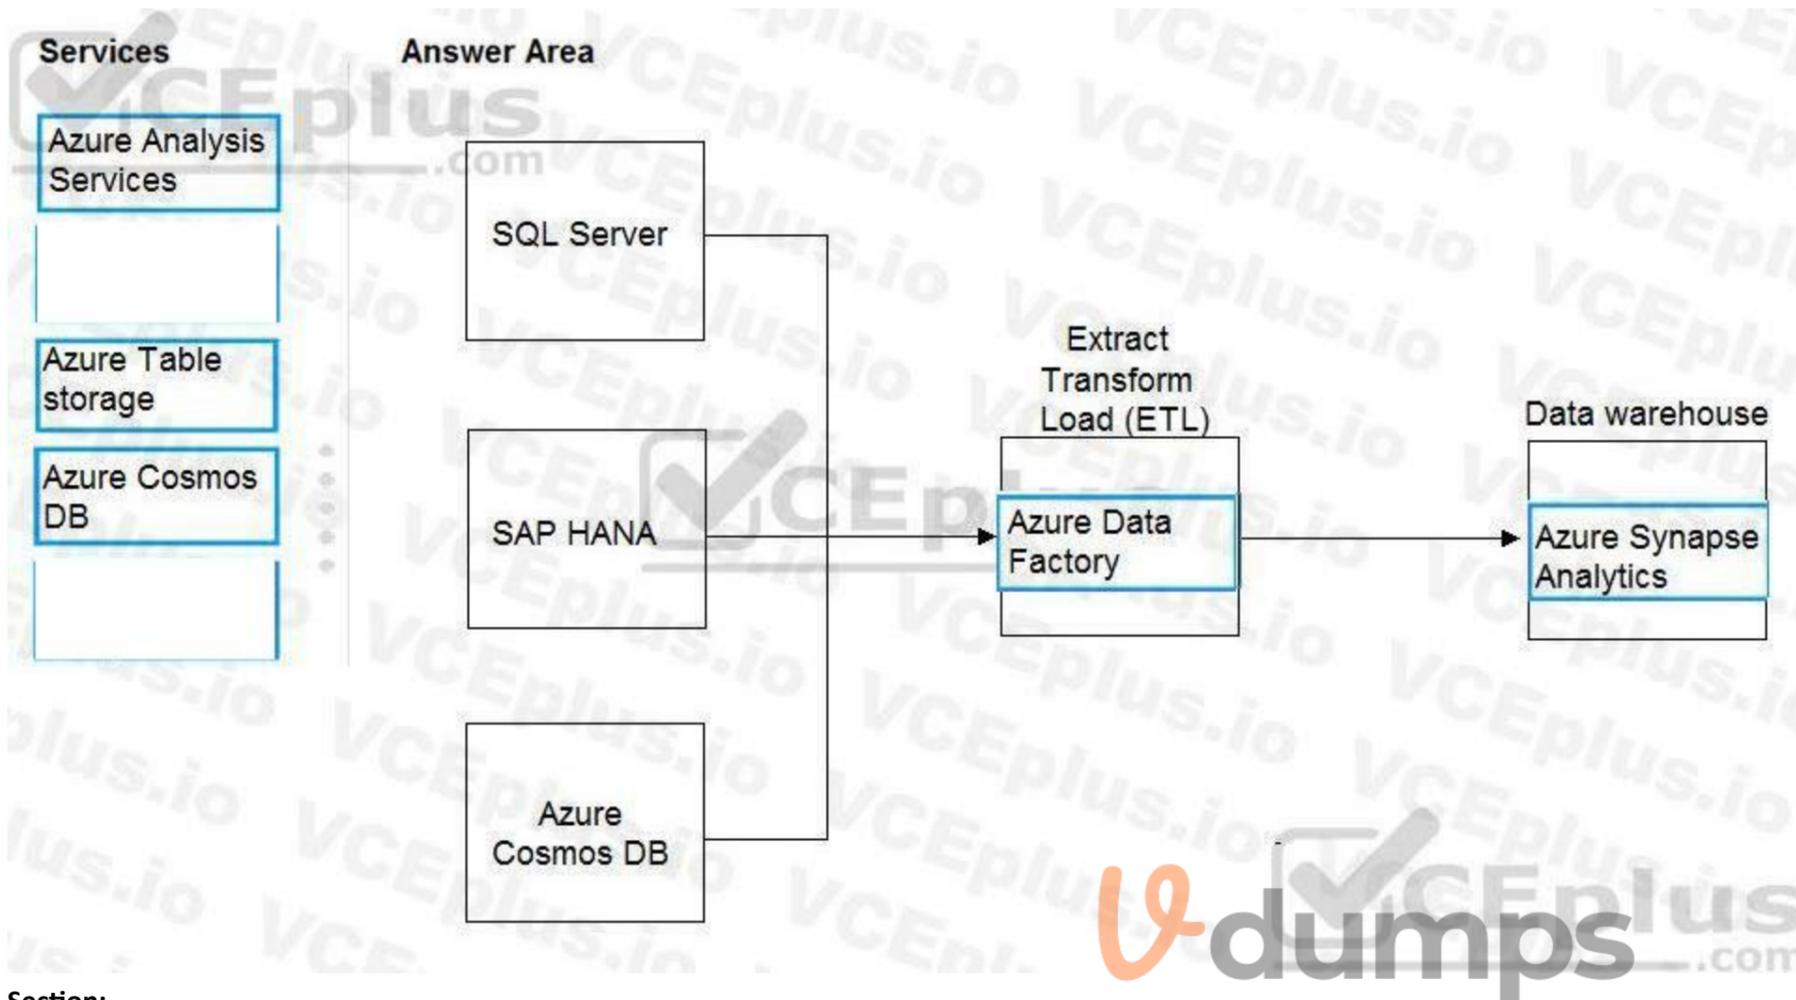
Service

Data warehouse

Service



Correct Answer:



Section:

Explanation:

Box 1: Azure Data factory

Relevant Azure service for the three ETL phases are Azure Data Factory and SQL Server Integration Services (SSIS).

Box 2: Azure Synapse Analytics

You can copy and transform data in Azure Synapse Analytics by using Azure Data Factory

Note: Azure Synapse Analytics connector is supported for the following activities:

Copy activity with supported source/sink matrix table

Mapping data flow

Lookup activity

GetMetadata activity

Reference:

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

<https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-sql-data-warehouse>

QUESTION 40

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:

Data Types	Answer Area
Image files	Data type Azure Blob storage
Key/value pairs	Data type Azure Cosmos DB Gremlin API
Relationships between employees	Data type Azure Table storage

Correct Answer:

Data Types	Answer Area
	Image files Azure Blob storage
	Key/value pairs Azure Cosmos DB Gremlin API
	Relationships between employees Azure Table storage

Section:

Explanation:

Box 1: Image files

Azure Blob storage is suitable for image files.

Box 2:Key/value pairs

Azure CosmosDB table API is a key-value storage hosted in the cloud.

Box 3: Relationship between employees

One-to-many relationships between business domain objects occur frequently: for example, one department has many employees. There are several ways to implement one-to-many relationships in the Azure Table service.

Reference:

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

QUESTION 41

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.	<input type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure provide configurable scaling options.	<input type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure reduce the administrative overhead for managing hardware.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
Platform as a service (PaaS) database offerings in Azure provide built-in high availability.	<input checked="" type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure provide configurable scaling options.	<input checked="" type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure reduce the administrative overhead for managing hardware.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

Reference:

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

QUESTION 42

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

dbo.Products: ▼

- A column
- A database
- A table
- An index

ProductName: ▼

- A column
- A database
- A table
- An index

Answer Area:

Answer Area

dbo.Products: ▼

- A column
- A database
- A table
- An index

ProductName: ▼

- A column
- A database
- A table
- An index

Section:

Explanation:

QUESTION 43

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 patches to Azure SQL databases regularly.	<input type="radio"/>	<input type="radio"/>
You need a Microsoft 365 subscription to create an Azure SQL database.	<input type="radio"/>	<input type="radio"/>
You can use existing Microsoft SQL Server licenses to reduce the cost of Azure SQL databases.	<input checked="" type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
You must apply patches to Azure SQL databases regularly.	<input type="radio"/>	<input checked="" type="radio"/>
You need a Microsoft 365 subscription to create an Azure SQL database.	<input type="radio"/>	<input checked="" type="radio"/>
You can use existing Microsoft SQL Server licenses to reduce the cost of Azure SQL databases.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

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 44

HOTSPOT

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

Hot Area:

Answer Area

Transparent Data Encryption (TDE) encrypts

- a column to protect data at rest and in transit.
- queries and their results in order to protect data in transit.
- the database to protect data at rest.
- the server to protect data at rest.

Answer Area:

Answer Area

Transparent Data Encryption (TDE) encrypts

- a column to protect data at rest and in transit.
- queries and their results in order to protect data in transit.
- the database to protect data at rest.
- the server to protect data at rest.



Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-tde-overview?tabs=azure-portal>

QUESTION 45

HOTSPOT

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

Hot Area:

Answer Area

By default, each Azure SQL database is protected by

- a network security group (NSG).
- a server-level firewall.
- Azure Firewall.
- Azure Front Door.

Answer Area:

Answer Area

By default, each Azure SQL database is protected by

- a network security group (NSG).
- a server-level firewall.
- Azure Firewall.
- Azure Front Door.

Section:

Explanation:

When you create a new server in Azure SQL Database or Azure Synapse Analytics named mysqlserver, for example, a server-level firewall blocks all access to the public endpoint for the server

Reference:

<https://docs.microsoft.com/en-us/azure/security/fundamentals/infrastructure-sql>

QUESTION 46

HOTSPOT

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

Hot Area:



Answer Area

A relational database must be used when

- a dynamic schema is required.
- data will be stored as key/value pairs.
- storing large images and videos.
- strong consistency guarantees are required.

Answer Area:

Answer Area

A relational database must be used when

- a dynamic schema is required.
- data will be stored as key/value pairs.
- storing large images and videos.
- strong consistency guarantees are required.

Section:

Explanation:

QUESTION 47

HOTSPOT

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

Hot Area:

Answer Area

Relational data uses

▼
collections
columns
keys
partitions

 to enforce relationships between different tables.

Answer Area:

Answer Area

Relational data uses

▼
collections
columns
keys
partitions

 to enforce relationships between different tables.

Section:

Explanation:

Reference:

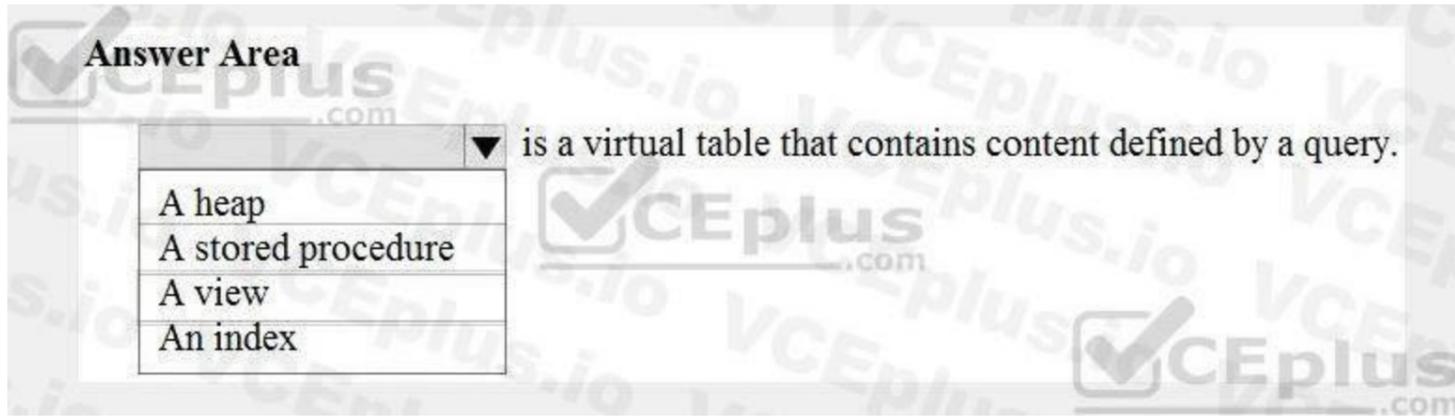
<https://teachcomputerscience.com/relational-databases/>

QUESTION 48

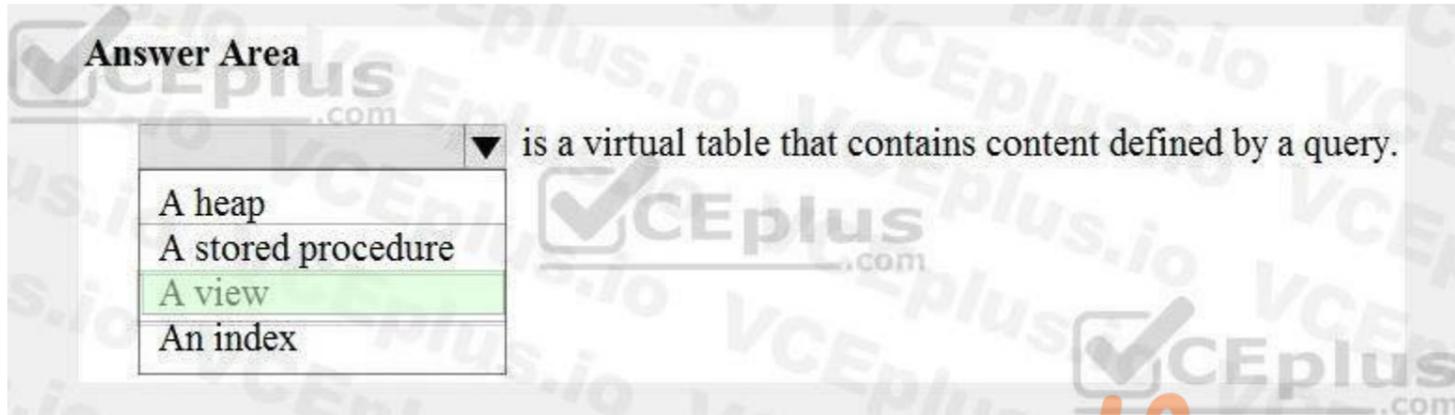
HOTSPOT

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

Hot Area:



Answer Area:



Section:

Explanation:

Reference:

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

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:

 **dumps**

Answer Area

Statements

	Yes	No
Relational database tables contain columns and rows	<input type="radio"/>	<input type="radio"/>
Indexes in a relational database describe the data types in a table	<input type="radio"/>	<input type="radio"/>
A database view is a virtual table whose content is defined by a query	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements

	Yes	No
Relational database tables contain columns and rows	<input checked="" type="radio"/>	<input type="radio"/>
Indexes in a relational database describe the data types in a table	<input type="radio"/>	<input checked="" type="radio"/>
A database view is a virtual table whose content is defined by a query	<input checked="" type="radio"/>	<input type="radio"/>

Section:

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 50

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
Copying data to Azure Data Lake Storage from another Azure region results in extra cost.	<input type="radio"/>	<input type="radio"/>
You can use blob, table, and file storage in the same Azure Storage account.	<input type="radio"/>	<input type="radio"/>
You implement Azure Data Lake Storage by creating an Azure Storage account.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Statements	Yes	No
Copying data to Azure Data Lake Storage from another Azure region results in extra cost.	<input checked="" type="radio"/>	<input type="radio"/>
You can use blob, table, and file storage in the same Azure Storage account.	<input checked="" type="radio"/>	<input type="radio"/>
You implement Azure Data Lake Storage by creating an Azure Storage account.	<input checked="" type="radio"/>	<input type="radio"/>



Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-get-started-portal>

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-overview>

<https://azure.microsoft.com/en-us/pricing/details/bandwidth/>

QUESTION 51

HOTSPOT

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

Hot Area:

Answer Area

When using the Azure Cosmos DB Gremlin API, the container resource type is projected as a

graph.
table.
partition key.
document.

Answer Area:

Answer Area

When using the Azure Cosmos DB Gremlin API, the container resource type is projected as a

graph.
table.
partition key.
document.

Vdumps

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/create-graph-gremlin-console>

QUESTION 52

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
Azure Table storage supports multiple read replicas.	<input type="radio"/>	<input type="radio"/>
Azure Table storage supports multiple write regions.	<input type="radio"/>	<input type="radio"/>
The Azure Cosmos DB Table API supports multiple read replicas.	<input type="radio"/>	<input type="radio"/>
The Azure Cosmos DB Table API supports multiple write regions.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Statements	Yes	No
Azure Table storage supports multiple read replicas.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Table storage supports multiple write regions.	<input type="radio"/>	<input checked="" type="radio"/>
The Azure Cosmos DB Table API supports multiple read replicas.	<input checked="" type="radio"/>	<input type="radio"/>
The Azure Cosmos DB Table API supports multiple write regions.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/high-availability>

QUESTION 53

HOTSPOT

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

Hot Area:

Answer Area

To configure an Azure Storage account to support both security at the folder level and atomic directory manipulation,

- enable the hierarchical namespace.
- set Account kind to BlobStorage.
- set Performance to Premium.
- set Replication to Read-access geo-redundant storage (RA-GRS).

Answer Area:

Answer Area

To configure an Azure Storage account to support both security at the folder level and atomic directory manipulation,

- enable the hierarchical namespace.
- set Account kind to BlobStorage.
- set Performance to Premium.
- set Replication to Read-access geo-redundant storage (RA-GRS).

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-namespace>

QUESTION 54

HOTSPOT

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

Hot Area:

Answer Area

Relational data uses

	▼
collections	
rows	
keys	
partitions	

to enforce relationships between different tables.

Answer Area:

Answer Area

Relational data uses

	▼
collections	
rows	
keys	
partitions	

to enforce relationships between different tables.



Section:

Explanation:

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

QUESTION 55

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
Stream processing has access to the most recent data received or data within a rolling time window.	<input type="radio"/>	<input type="radio"/>
Batch processing must occur immediately and have latency in the order of seconds or milliseconds.	<input type="radio"/>	<input type="radio"/>
Stream processing is used for simple response functions, aggregates, or calculations such as rolling averages.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
Stream processing has access to the most recent data received or data within a rolling time window.	<input checked="" type="radio"/>	<input type="radio"/>
Batch processing must occur immediately and have latency in the order of seconds or milliseconds.	<input type="radio"/>	<input checked="" type="radio"/>
Stream processing is used for simple response functions, aggregates, or calculations such as rolling averages.	<input type="radio"/>	<input checked="" type="radio"/>

Section:

Explanation:

Reference:

<https://www.precisely.com/blog/big-data/big-data-101-batch-stream-processing>

QUESTION 56

HOTSPOT

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

Hot Area:

Answer Area

	▼	is an Azure Synapse Analytics feature that retrieves data from Azure Blob storage by using Transact-SQL queries.
A hierarchical namespace		
Azure Synapse Link		
PolyBase		
The Data Movement Service (DMS)		

Answer Area:

Answer Area

	▼	is an Azure Synapse Analytics feature that retrieves data from Azure Blob storage by using Transact-SQL queries.
A hierarchical namespace		
Azure Synapse Link		
PolyBase		
The Data Movement Service (DMS)		



Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-guide?view=sql-server-ver15>

QUESTION 57

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

- line chart
- matrix
- scatter
- treemap

the relative value of each item. The chart can be hierarchical, displaying data as a set of rectangles nested within the main rectangle.

Answer Area:

Answer Area

A is a chart of colored rectangles. The size of the rectangles represent

- line chart
- matrix
- scatter
- treemap

the relative value of each item. The chart can be hierarchical, displaying data as a set of rectangles nested within the main rectangle.



Section:

Explanation:

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

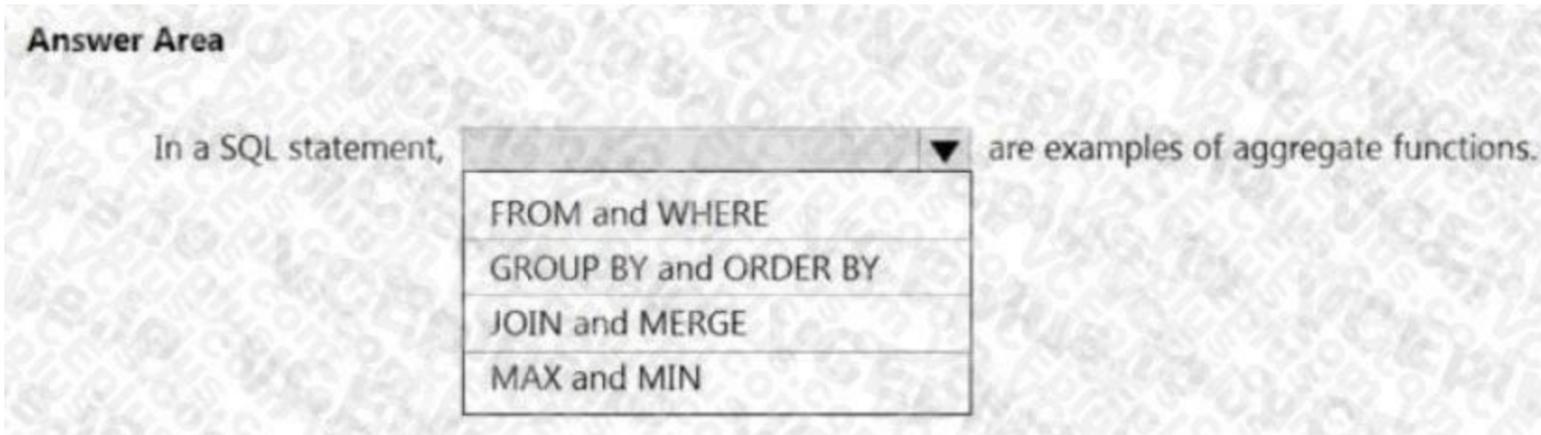
Reference: <https://docs.microsoft.com/en-us/sql/reporting-services/report-design/tree-map-and-sunburst-charts-in-reportingservices?view=sql-server-ver15>

QUESTION 58

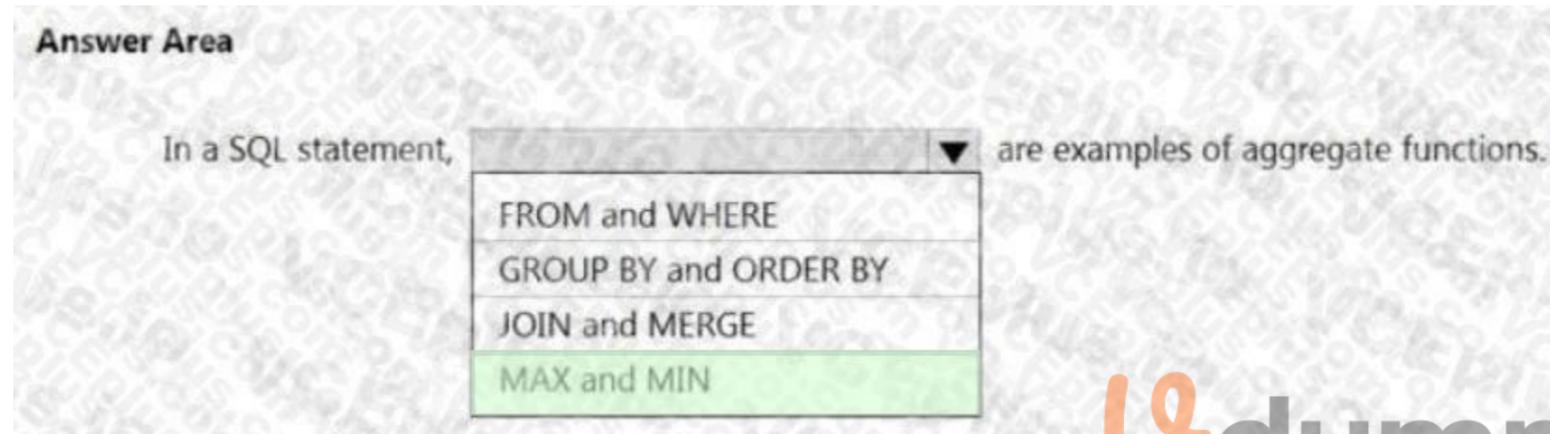
HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



Answer Area:



Section:

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-ver15>

QUESTION 59

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.

 **vdumps**

Hot Area:

Answer Area	Statements	Yes	No
	A job that processes sales data once daily is an example of a batch workload.	<input type="checkbox"/>	<input type="checkbox"/>
	A job that calculates a rolling average temperature reading is an example of a streaming workload.	<input type="checkbox"/>	<input type="checkbox"/>
	A job that calculates average revenue per product for the last month is an example of a streaming workload.	<input type="checkbox"/>	<input type="checkbox"/>

Answer Area:

Answer Area	Statements	Yes	No
	A job that processes sales data once daily is an example of a batch workload.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
	A job that calculates a rolling average temperature reading is an example of a streaming workload.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
	A job that calculates average revenue per product for the last month is an example of a streaming workload.	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Section:

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 60

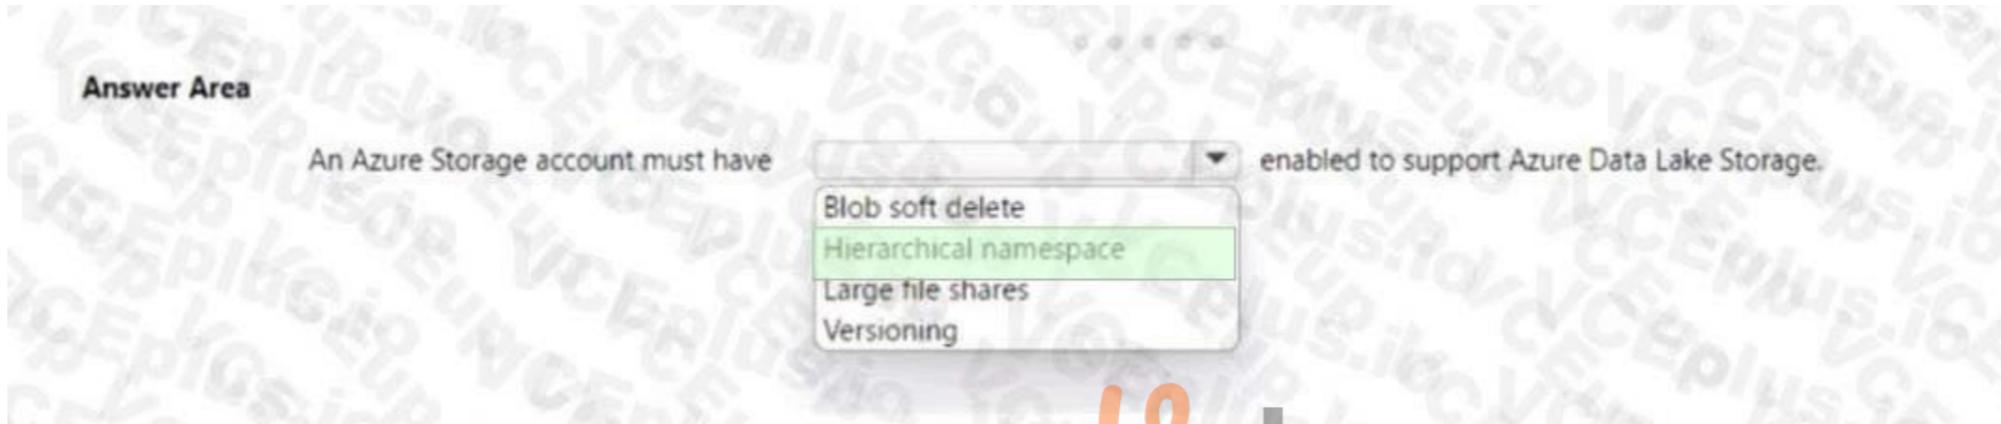
HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



Answer Area:



Section:

Explanation:

<https://learn.microsoft.com/en-us/dynamics365/customer-insights/own-data-lake-storage>

QUESTION 61

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



Answer Area:

Answer Area

Microsoft Power BI Desktop provides a full range of Power BI data modeling and report editing features.

- Microsoft Power BI Desktop
- The Microsoft Power BI service
- The Microsoft Power BI phone app

provides a full range of Power BI data modeling and report editing features.

Section:

Explanation:

QUESTION 62

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 Area:

Answer Area	Statements	Yes	No
	Normalizing a database increases the throughput of writing transactions.	<input checked="" type="radio"/>	<input type="radio"/>
	Analytics systems are more normalized than transactional systems.	<input type="radio"/>	<input checked="" type="radio"/>
	Normalizing a database results in queries that require more joins.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

QUESTION 63

HOTSPOT

Select the answer that correctly completes the sentence.



Hot Area:

Answer Area

Microsoft Power BI Desktop provides a full range of Power BI data modeling and report editing features.

- Microsoft Power BI Desktop
- The Microsoft Power BI service
- The Microsoft Power BI phone app

Answer Area:

Answer Area

Microsoft Power BI Desktop provides a full range of Power BI data modeling and report editing features.

- Microsoft Power BI Desktop
- The Microsoft Power BI service
- The Microsoft Power BI phone app

Section:

Explanation:

QUESTION 64

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

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

- data that is fully processed before being loaded to the target data store.
- 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.

Content Display

Answer Area:

Answer Area

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

- data that is fully processed before being loaded to the target data store.
- 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.

Content Display

Section:

Explanation:



QUESTION 65

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

In Microsoft Power BI, a

- dashboard
- dashboard
- dataset
- report
- tile

is a collection of visualizations that appear together on one or more pages.

Answer Area:

Answer Area

In Microsoft Power BI, a

- dashboard
- dashboard
- dataset
- report
- tile

is a collection of visualizations that appear together on one or more pages.

Section:

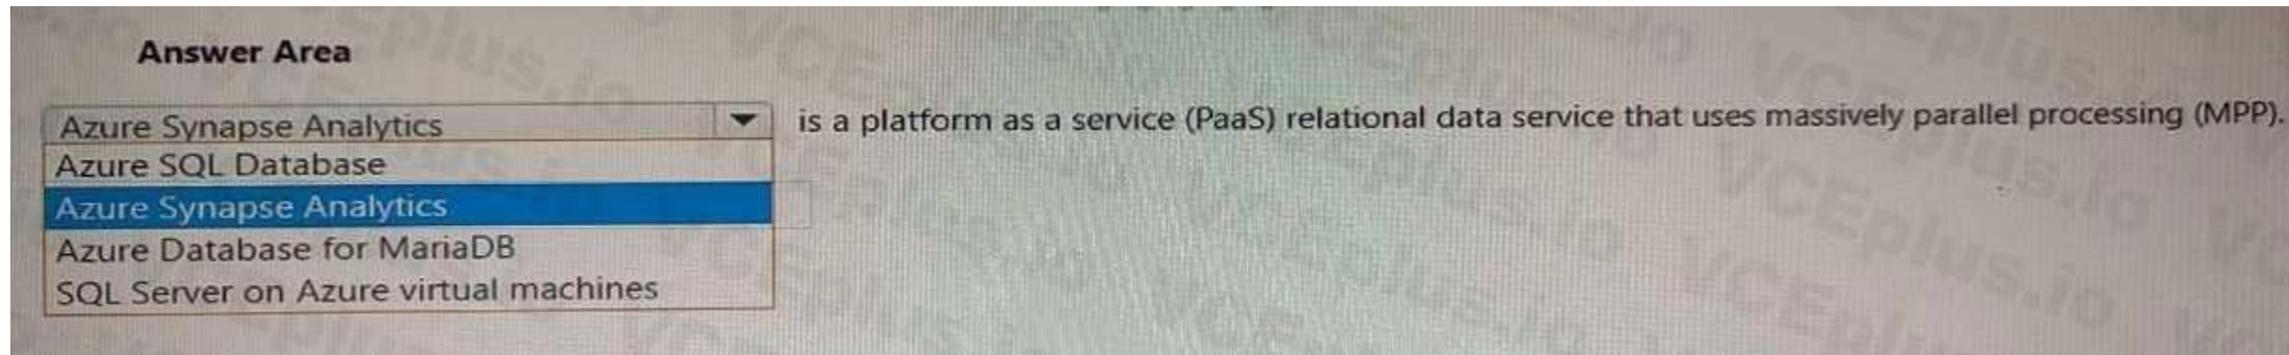
Explanation:

QUESTION 66

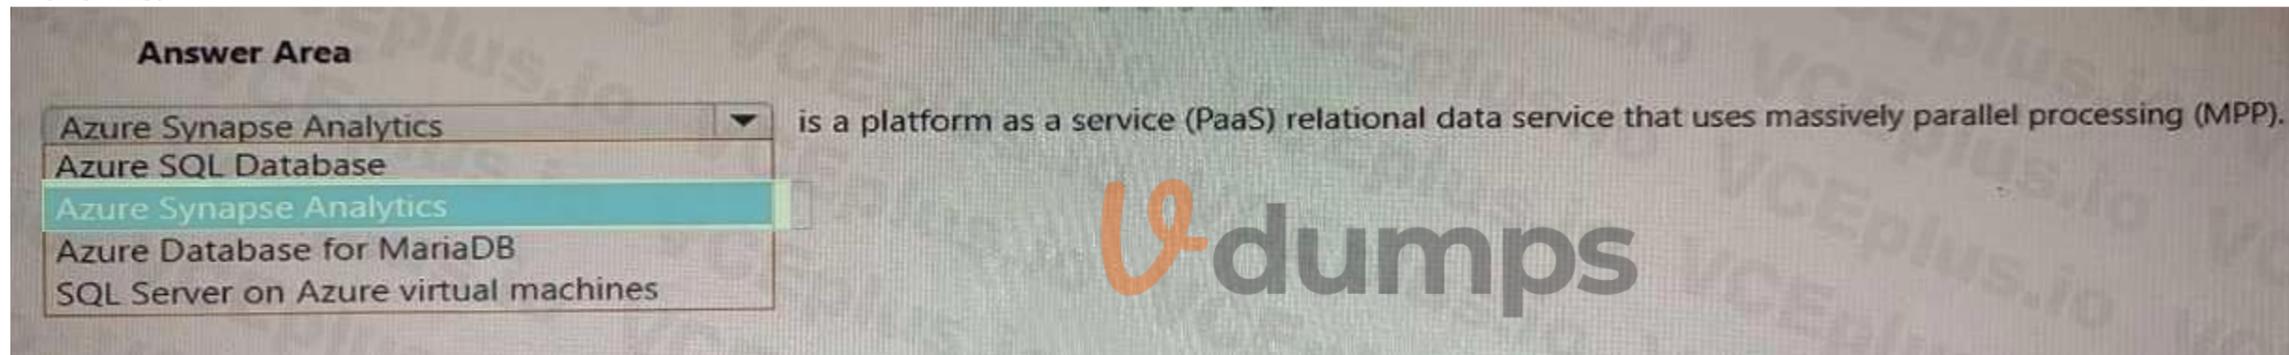
HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



Answer Area:



Section:

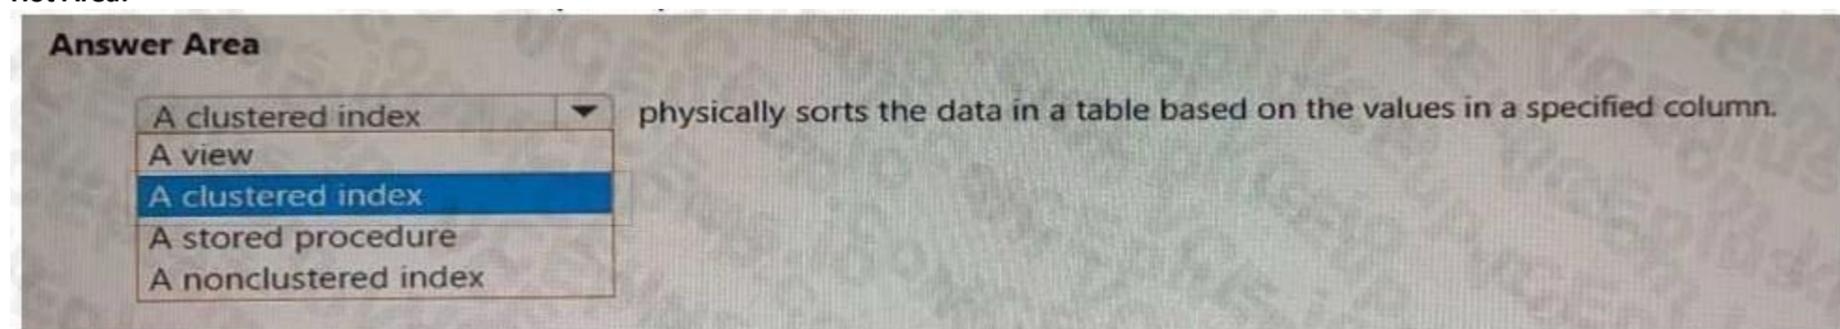
Explanation:

QUESTION 67

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



Answer Area:

Answer Area

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

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

Section:

Explanation:

QUESTION 68

HOTSPOT

Select the answer that correctly completes the sentence. 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.	<input type="radio"/>	<input type="radio"/>
You need a Microsoft 365 subscription to create an Azure SQL database.	<input type="radio"/>	<input type="radio"/>
You can use existing Microsoft SQL Server licenses to reduce the cost of Azure SQL databases.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
You must apply operating system updates to Azure SQL databases regularly.	<input type="radio"/>	<input checked="" type="radio"/>
You need a Microsoft 365 subscription to create an Azure SQL database.	<input type="radio"/>	<input checked="" type="radio"/>
You can use existing Microsoft SQL Server licenses to reduce the cost of Azure SQL databases.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

QUESTION 69

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

In a relational database, the uniquely identifies each row in a table.

- primary key
- foreign key
- primary key
- WHERE clause
- nonclustered index

Answer Area:

Answer Area

In a relational database, the uniquely identifies each row in a table.

- primary key
- foreign key
- primary key
- WHERE clause
- nonclustered index

Section:

Explanation:



QUESTION 70

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

presents content defined by a query.

- A view
- A heap
- A stored procedure
- A view
- An index

Answer Area:

Select the answer that correctly completes the sentence.

Answer Area

A view presents content defined by a query.

- A view
- A heap
- A stored procedure
- A view
- An index

Section:

Explanation:

QUESTION 71

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



Answer Area

Column family databases natively support the analysis of relationships between entities.

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

Answer Area:

Answer Area

Column family databases natively support the analysis of relationships between entities.

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

Section:

Explanation:

QUESTION 72

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A data engineer is responsible for creating visuals and charts that help a company make informed decisions.

- A data engineer
- A data analyst
- A data engineer
- A data scientist
- A database administrator

Answer Area:

Answer Area

A data engineer is responsible for creating visuals and charts that help a company make informed decisions.

- A data engineer
- A data analyst
- A data engineer
- A data scientist
- A database administrator

Section:

Explanation:

QUESTION 73

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

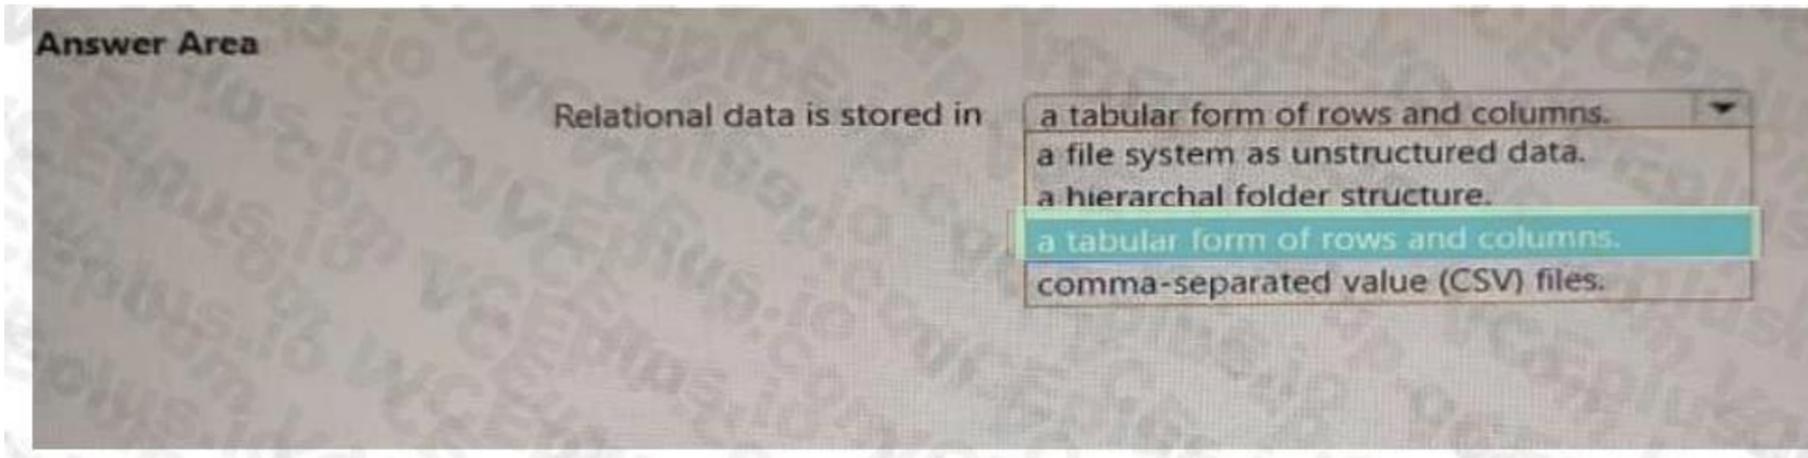
Answer Area

Relational data is stored in

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

Answer Area:





Section:

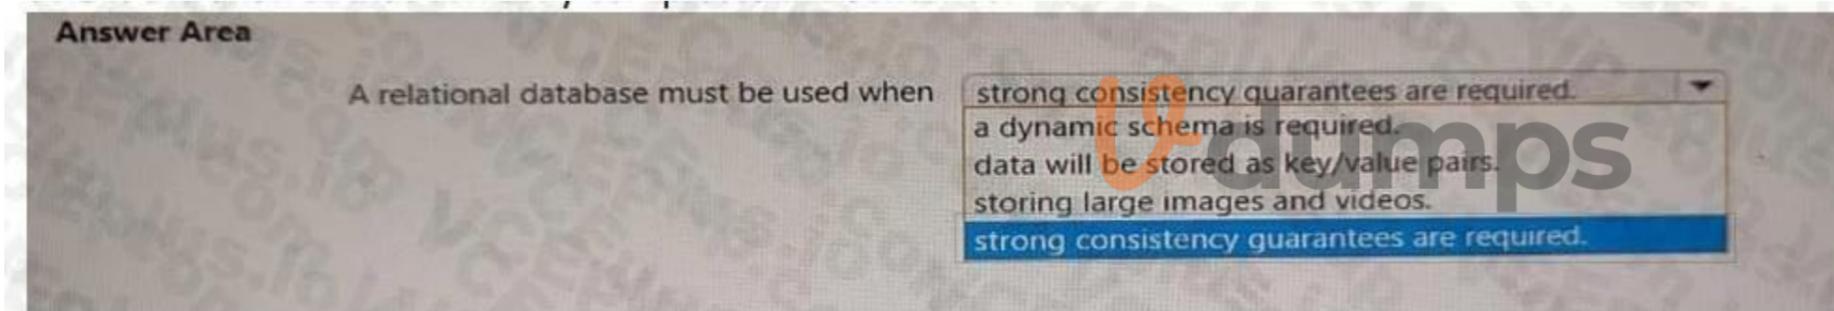
Explanation:

QUESTION 74

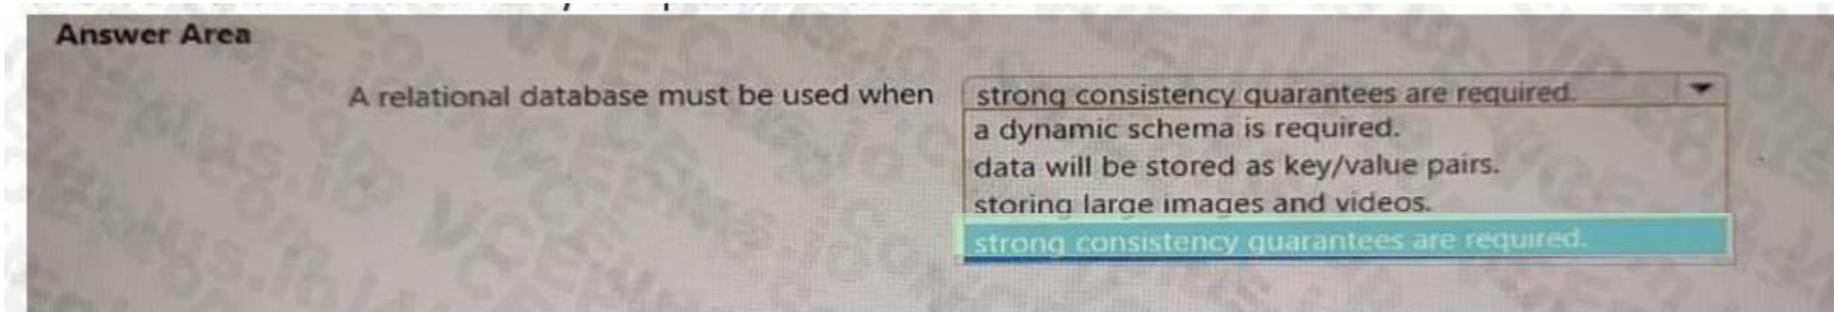
HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



Answer Area:



Section:

Explanation:

QUESTION 75

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

Backing up and restoring data are the responsibilities of a

- data engineer.
- data analyst.
- data engineer.**
- data scientist.
- database administrator.

Answer Area:

Answer Area

Backing up and restoring data are the responsibilities of a

- data engineer.
- data analyst.
- data engineer.**
- data scientist.
- database administrator.

Section:

Explanation:

QUESTION 76

HOTSPOT

Select the answer that correctly completes the sentence.



Hot Area:

Answer Area

A key/value data store is optimized for

- simple lookups.
- enforcing constraints.
- simple lookups.**
- table joins.

Answer Area:

Answer Area

A key/value data store is optimized for

- simple lookups.
- enforcing constraints.
- simple lookups.**
- table joins.

Section:

Explanation:

QUESTION 77

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

Online analytical processing (OLAP) systems are particularly suitable for analytical workloads because data is

- optimized for transactions.
- limited in size.
- optimized for transactions.
- pre-aggregated.
- read-only.

Answer Area:

Answer Area

Online analytical processing (OLAP) systems are particularly suitable for analytical workloads because data is

- optimized for transactions.
- limited in size.
- optimized for transactions.
- pre-aggregated.
- read-only.

Section:

Explanation:



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
You create Azure File Storage in a storage account.	<input checked="" type="radio"/>	<input type="radio"/>
The Premium tier of Azure File Storage uses HDD-based hardware in a datacenter.	<input type="radio"/>	<input checked="" type="radio"/>
The Standard tier of Azure File Storage uses SSDs.	<input type="radio"/>	<input checked="" type="radio"/>

Answer Area:



Section:

Explanation:

QUESTION 79

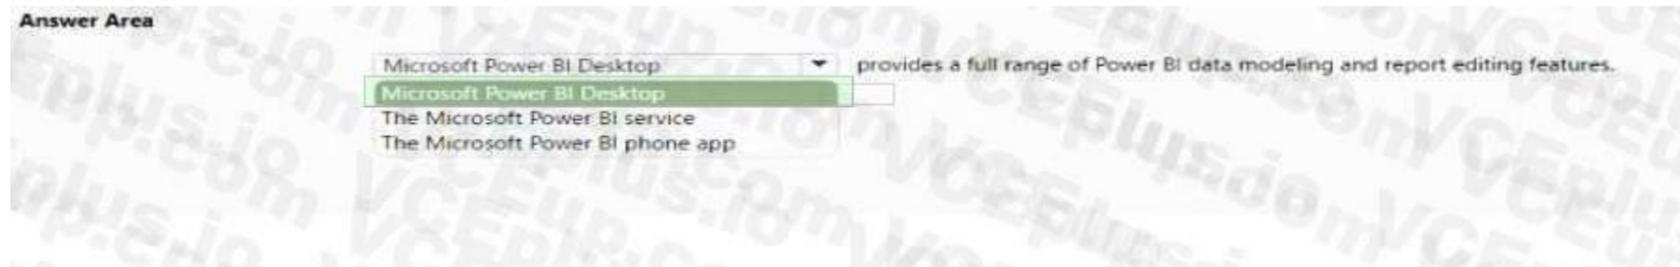
HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



Answer Area:



Section:

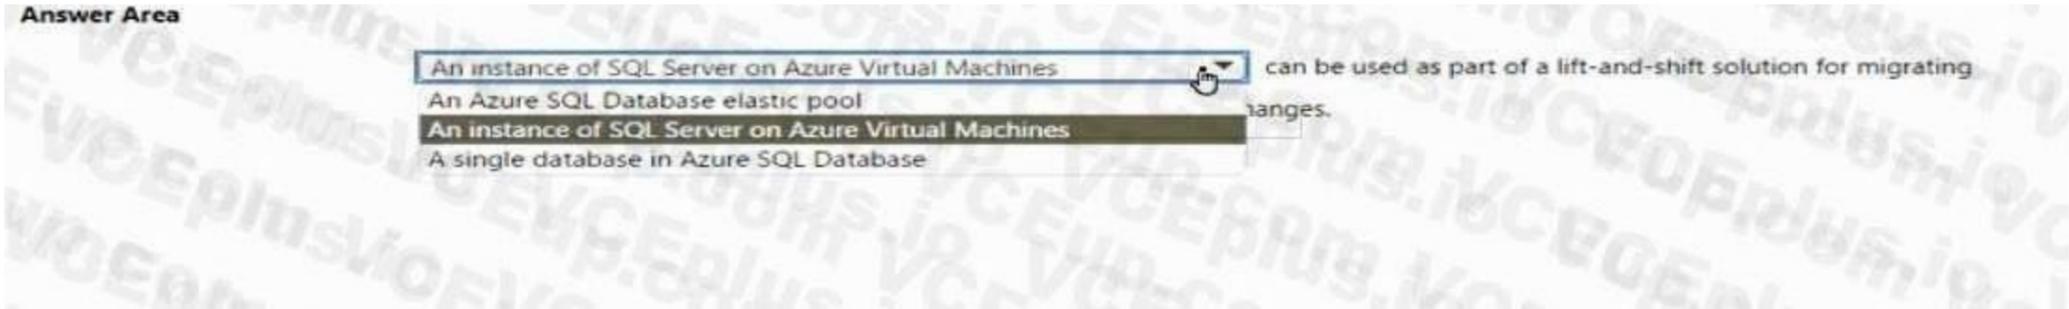
Explanation:

QUESTION 80

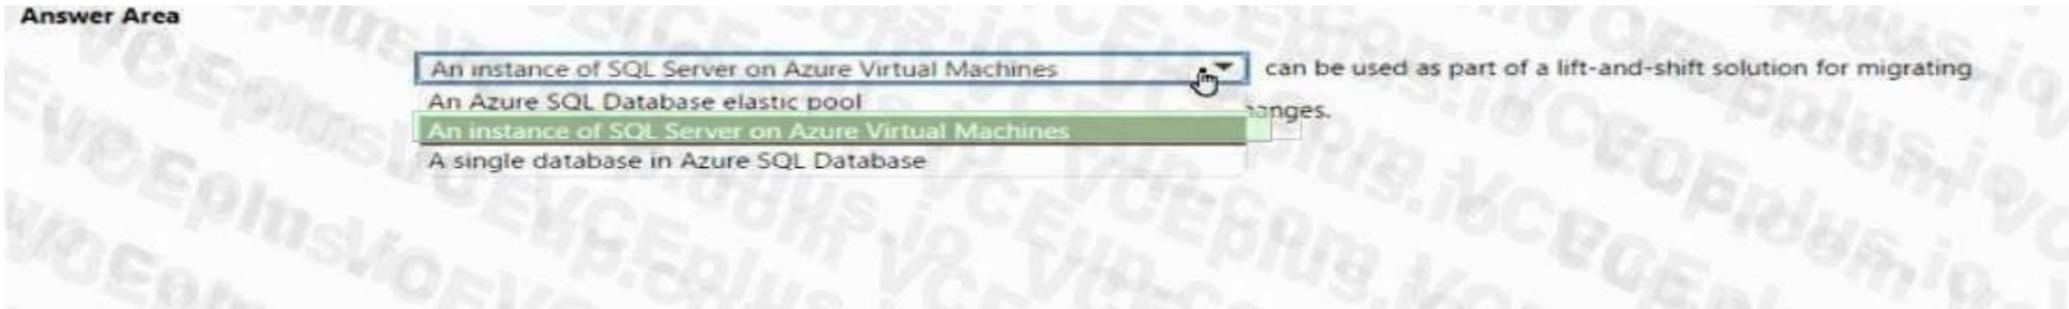
HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

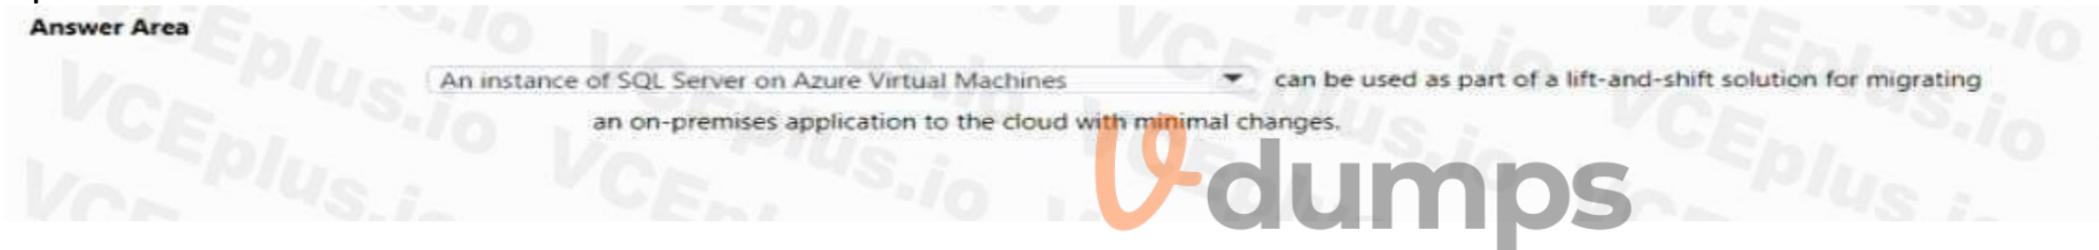


Answer Area:



Section:

Explanation:

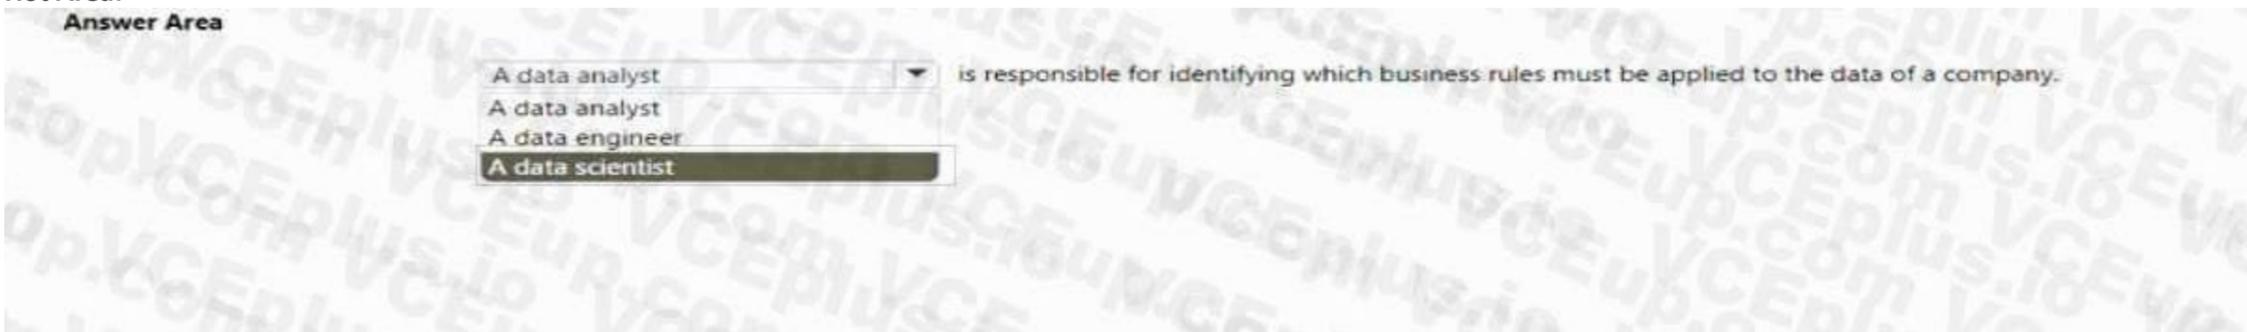


QUESTION 81

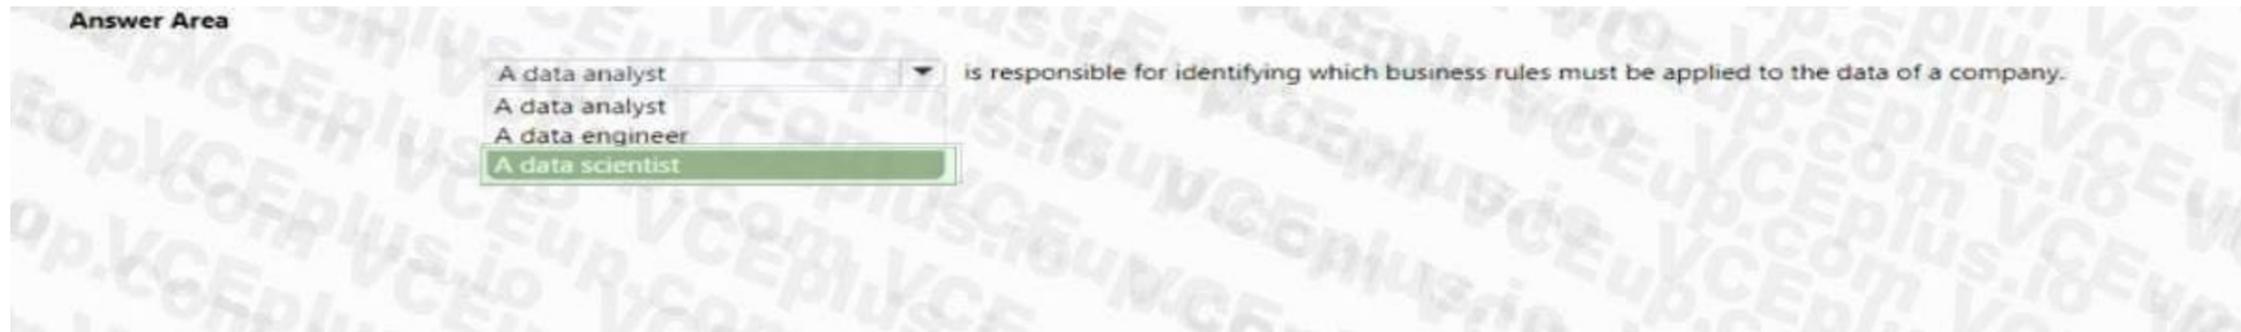
HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



Answer Area:



Section:

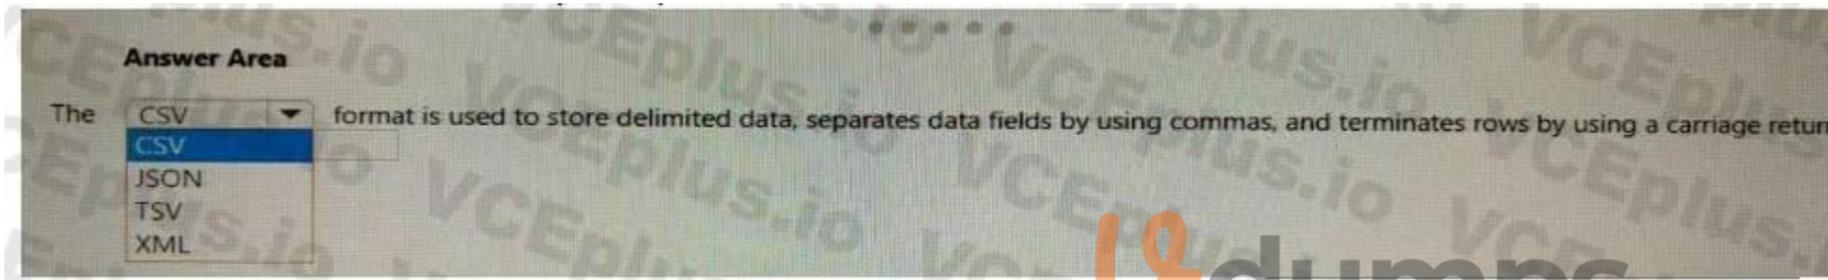
Explanation:

QUESTION 82

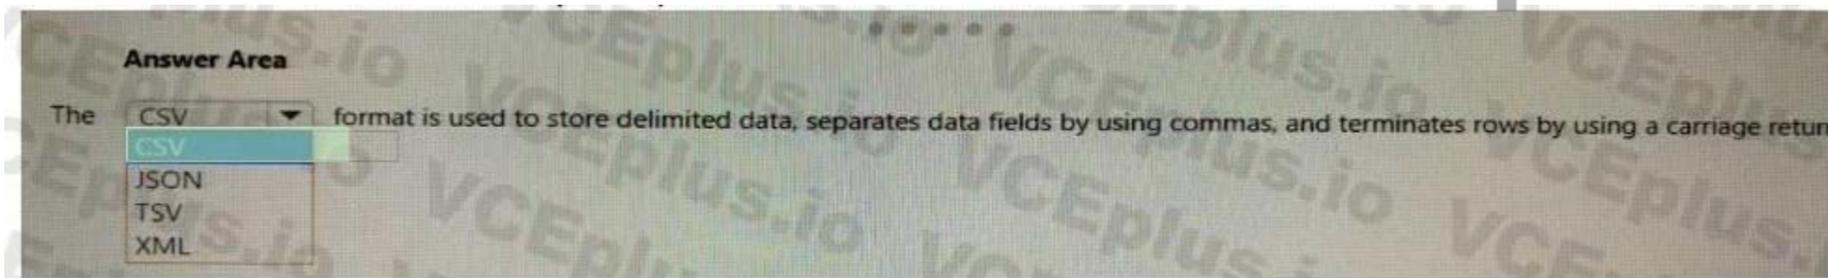
HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



Answer Area:



Section:

Explanation:

CSV

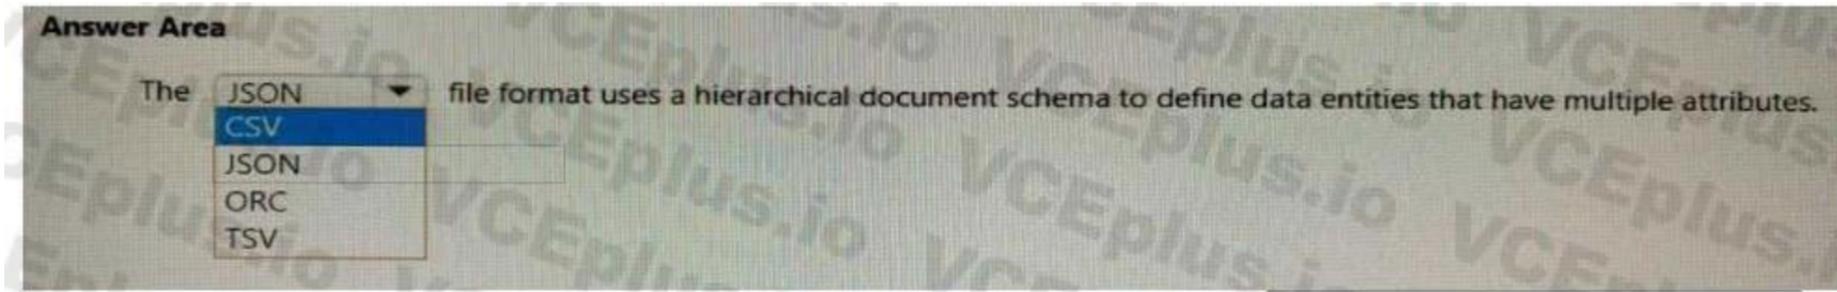
CSV (Comma-separated values): This is a text file format that uses commas (,) as column delimiters and carriage returns (\r\n) as row delimiters. CSV files are widely used for exchanging data between different applications and platforms¹²³.

QUESTION 83

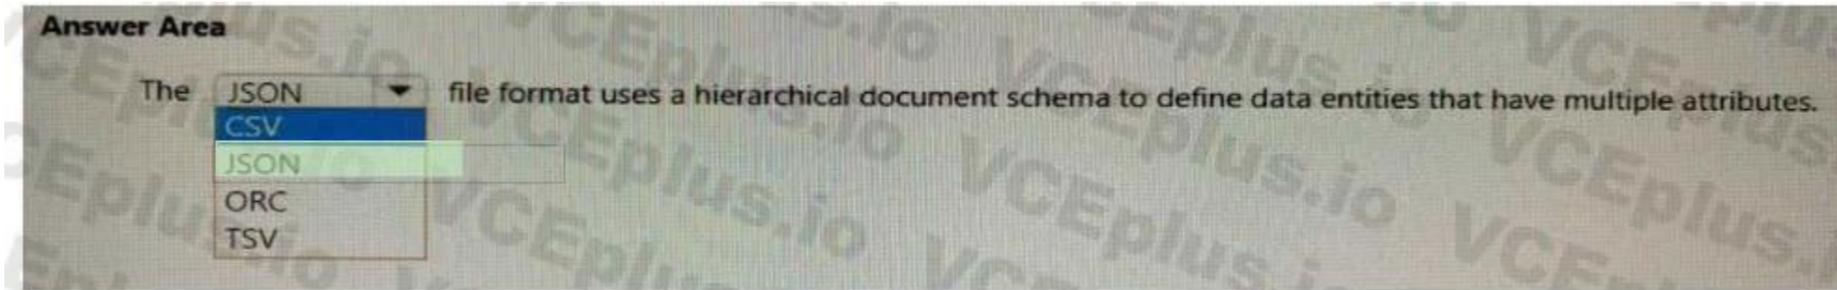
HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



Answer Area:



Section:

Explanation:

JSON (JavaScript Object Notation): This is a text file format that uses a human-readable and compact syntax to represent data as a collection of name-value pairs and arrays. JSON files can express complex and nested data structures by using objects and arrays as values¹²³.

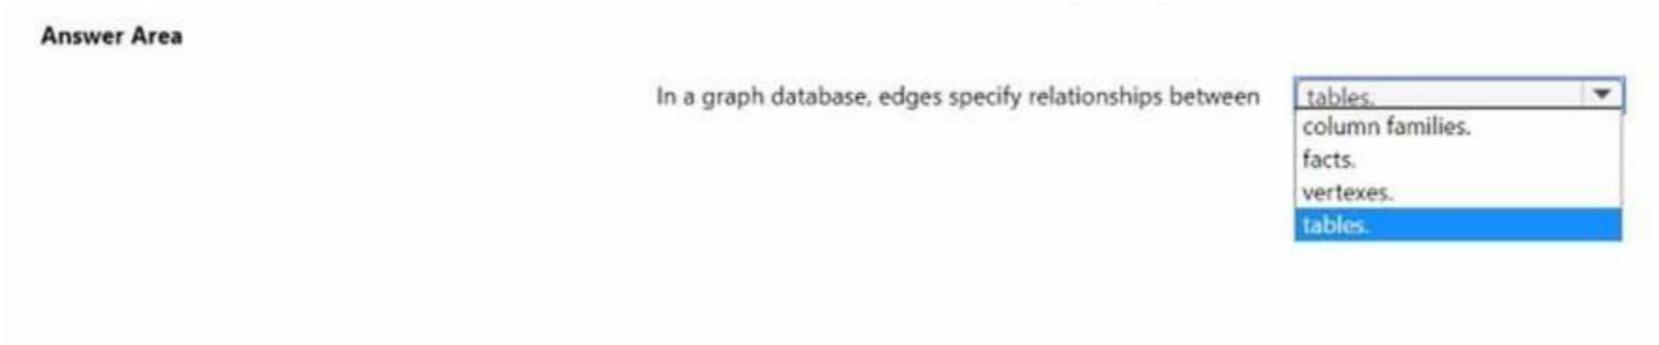
QUESTION 84

HOTSPOT

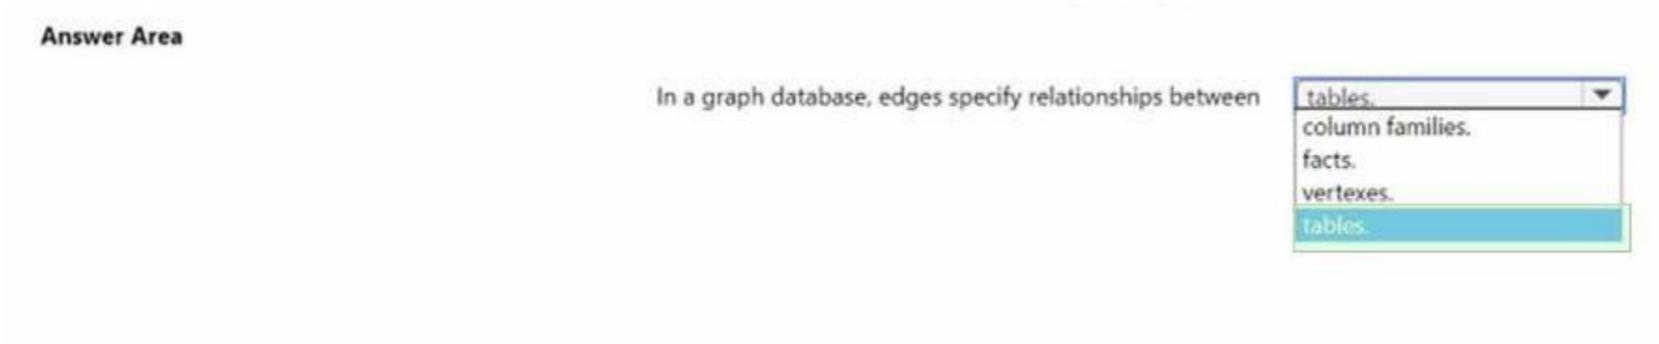
Select the answer that correctly completes the sentence.



Hot Area:



Answer Area:



Section:

Explanation:

QUESTION 85

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

The pipeline engine used in Azure Data Factory is also available in

- Azure Synapse Analytics.
- Azure Databricks.
- Azure HDInsight.
- Azure SQL Database.
- Azure Synapse Analytics.

Answer Area:

Answer Area

The pipeline engine used in Azure Data Factory is also available in

- Azure Synapse Analytics.
- Azure Databricks.
- Azure HDInsight.
- Azure SQL Database.
- Azure Synapse Analytics.

Section:

Explanation:

QUESTION 86

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

The Azure Cosmos DB enables the use of SELECT statements to retrieve documents from Azure Cosmos DB.

- for NoSQL
- for Apache Gremlin
- for MongoDB
- for Table

Answer Area:

Answer Area

The Azure Cosmos DB enables the use of SELECT statements to retrieve documents from Azure Cosmos DB.

Section:
Explanation:

QUESTION 87
HOTSPOT
Select the answer that correctly completes the sentence.

Hot Area:
Answer Area

A Microsoft Power BI enables users to create highly formatted, fixed-layout documents optimized for printing and archiving.

Answer Area:
Answer Area

A Microsoft Power BI enables users to create highly formatted, fixed-layout documents optimized for printing and archiving.

Section:
Explanation:

QUESTION 88
HOTSPOT
Select the answer that correctly completes the sentence.

Hot Area:
Answer Area

Batch workloads

Answer Area:

Answer Area

Batch workloads

- collect data and then process the data when a condition is met.
- always process data in memory, row-by-row.
- collect and process data at most once a day.
- process data as new data is received in near real-time.
- collect data and then process the data when a condition is met.

Section:

Explanation:

QUESTION 89

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. graph
- B. key/value
- C. document
- D. columnar

Correct Answer: A

Section:

Explanation:

Section: Describe how to work with non-relational data on Azure

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 90

What are two uses of data visualization? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Represent trends and patterns over time
- B. Implement machine learning to predict future values
- C. Communicate the significance of data
- D. Consistently implement business logic across reports

Correct Answer: A, C

Section:

Explanation:

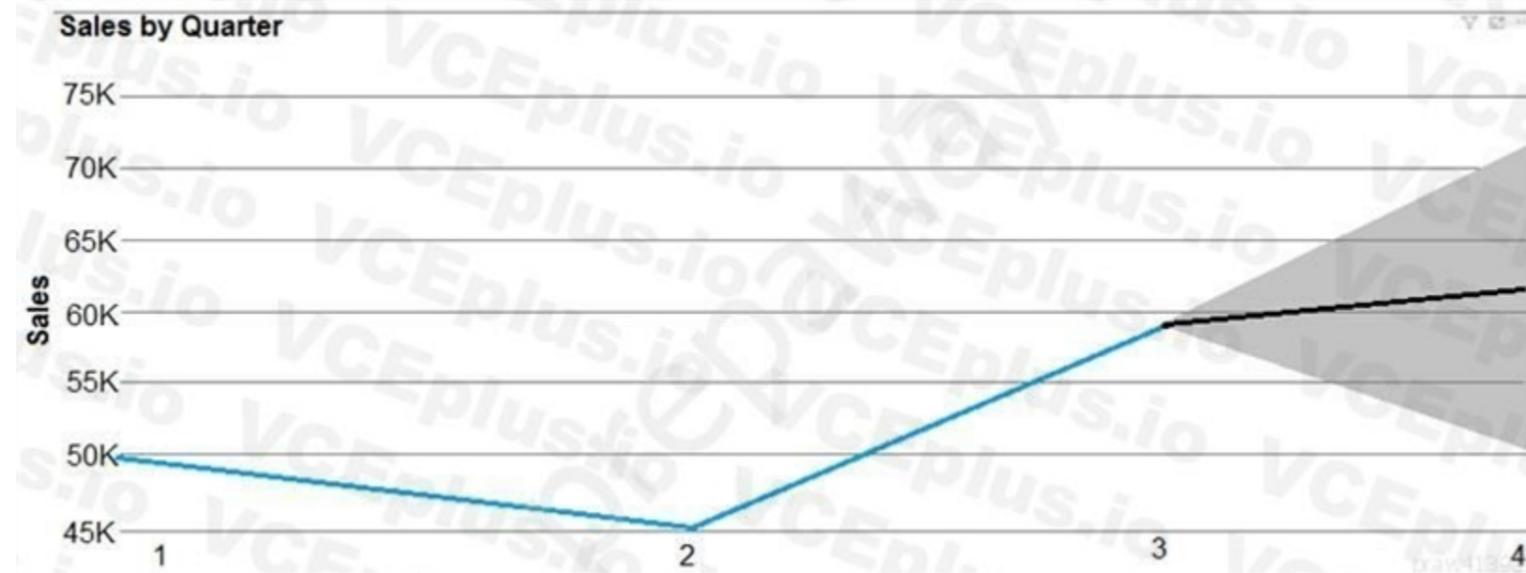
Section: Describe an analytics workload on Azure

Data visualization is a key component in being able to gain insight into your data. It helps make big and small data easier for humans to understand. It also makes it easier to detect patterns, trends, and outliers in groups of data. Data visualization brings data to help you find key business insights quickly and effectively.

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/spark/apache-spark-data-visualization>

QUESTION 91

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

Correct Answer: A

Section:

Explanation:

Predictive, to answer the question: What will happen?



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 92

HOTSPOT

 **vdumps**

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

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

- a target data store powerful enough to transform data.
- a data pipeline that includes a transformation engine.
- 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.

Answer Area:

Answer Area

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

- a target data store powerful enough to transform data.
- a data pipeline that includes a transformation engine.
- 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.

Section:

Explanation:

QUESTION 93

HOTSPOT

Select the answer that correctly completes the sentence.



Hot Area:

Answer Area

Relational data uses

- keys
- collections
- rows
- keys
- partitions

to enforce relationships between different tables.

Answer Area:

Answer Area

Relational data uses

- keys
- collections
- rows
- keys
- partitions

to enforce relationships between different tables.

Section:

Explanation:

QUESTION 94

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

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

Answer Area:

Answer Area

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

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



Section:

Explanation:

QUESTION 95

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A JSON document that contains customer entities with a variable number of properties is an example of

semi-structured	▼
relational	
semi-structured	
structured	
unstructured	

data.

Answer Area:

Answer Area

A JSON document that contains customer entities with a variable number of properties is an example of

semi-structured	▼
relational	
semi-structured	
structured	
unstructured	

data.

Section:

Explanation:

QUESTION 96

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

The logo for 'Vdumps' features a stylized orange 'V' followed by the word 'dumps' in a grey, lowercase, sans-serif font.

In an analytical model, any point where the dimensions intersect represents

a cube.	▼
a cube.	
an aggregated dimension.	
an aggregated measure.	

Answer Area:

Answer Area

In an analytical model, any point where the dimensions intersect represents

a cube.	▼
a cube.	
an aggregated dimension.	
an aggregated measure.	

Section:

Explanation:

QUESTION 97

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A data analyst
A data analyst
A data engineer
A database administrator

is responsible for finding hidden patterns in data and creating graphical reports that deliver information in a useful and appealing format.

Answer Area:

Answer Area

A data analyst
A data analyst
A data engineer
A database administrator

is responsible for finding hidden patterns in data and creating graphical reports that deliver information in a useful and appealing format.

Section:

Explanation:

QUESTION 98

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
In non-relational databases, data is always stored in tables.	<input type="radio"/>	<input type="radio"/>
In non-relational databases, data is always stored by using a fixed schema.	<input type="radio"/>	<input type="radio"/>
In non-relational databases, two entities in the same collection can have different sets of fields.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
In non-relational databases, data is always stored in tables.	<input type="radio"/>	<input checked="" type="radio"/>
In non-relational databases, data is always stored by using a fixed schema.	<input type="radio"/>	<input checked="" type="radio"/>
In non-relational databases, two entities in the same collection can have different sets of fields.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

QUESTION 99

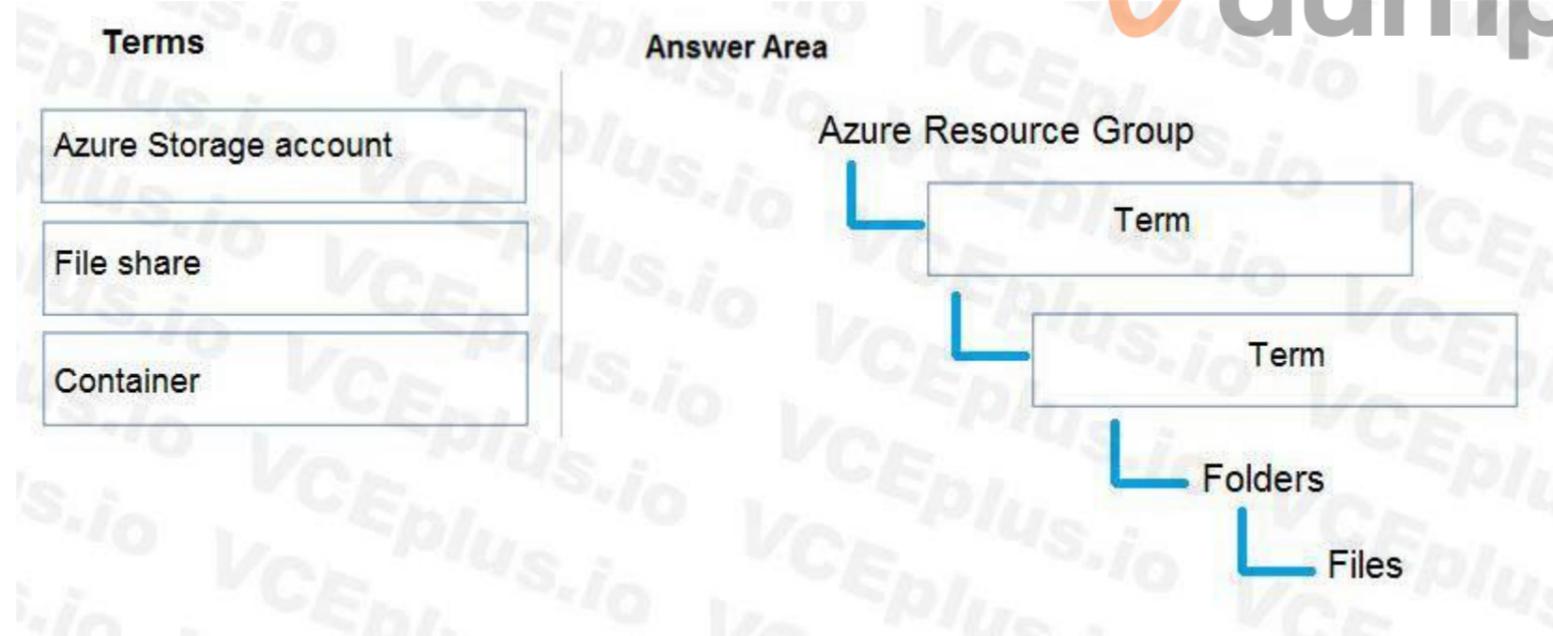
DRAG DROP

Match the Azure Data Lake Storage terms to the appropriate levels in the hierarchy.

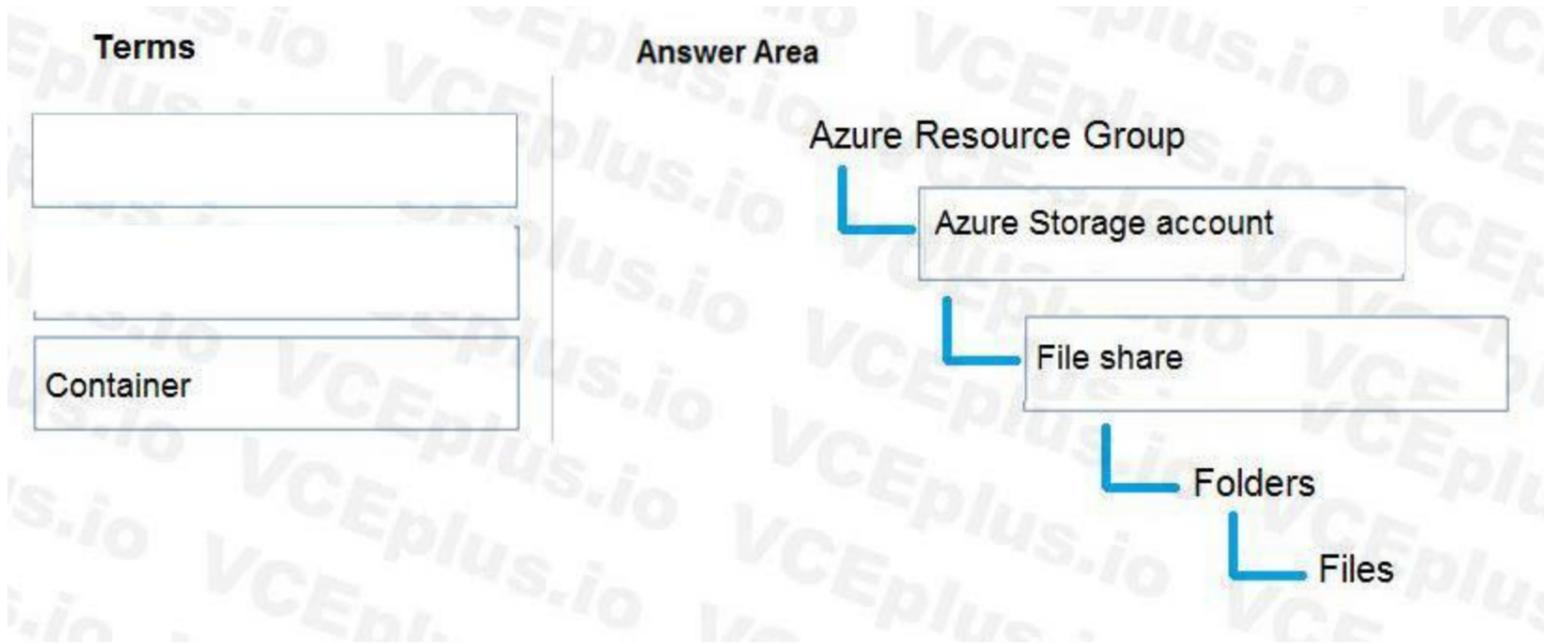
To answer, drag the appropriate term from the column on the left to its level on the right. Each term may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:



Correct Answer:



Section:

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: File share

Reference:

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



QUESTION 100

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:

Components	Answer Area
Dataset	Component A representation of data structures within data stores
Linked service	Component The information used to connect to external resources
Mapping data flow	Component A logical grouping of activities that performs a unit of work and can be scheduled
Pipeline	

Correct Answer:

Components	Answer Area						
	<table border="1"><tr><td>Dataset</td><td>A representation of data structures within data stores</td></tr><tr><td>Linked service</td><td>The information used to connect to external resources</td></tr><tr><td>Pipeline</td><td>A logical grouping of activities that performs a unit of work and can be scheduled</td></tr></table>	Dataset	A representation of data structures within data stores	Linked service	The information used to connect to external resources	Pipeline	A logical grouping of activities that performs a unit of work and can be scheduled
Dataset	A representation of data structures within data stores						
Linked service	The information used to connect to external resources						
Pipeline	A logical grouping of activities that performs a unit of work and can be scheduled						
Mapping data flow							

Section:

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.com/en-us/azure/data-factory/concepts-linked-services>

<https://docs.microsoft.com/en-us/azure/data-factory/concepts-pipelines-activities>



QUESTION 101

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:

Workload Types	Answer Area						
Batch	<table border="1"><tr><td>Workload type</td><td>Data for a product catalog will be loaded every 12 hours to a data warehouse.</td></tr><tr><td>Workload type</td><td>Data for online purchases will be loaded to a data warehouse as the purchases occur.</td></tr><tr><td>Workload type</td><td>Updates to inventory data will be loaded to a data warehouse every 1,000 transactions.</td></tr></table>	Workload type	Data for a product catalog will be loaded every 12 hours to a data warehouse.	Workload type	Data for online purchases will be loaded to a data warehouse as the purchases occur.	Workload type	Updates to inventory data will be loaded to a data warehouse every 1,000 transactions.
Workload type	Data for a product catalog will be loaded every 12 hours to a data warehouse.						
Workload type	Data for online purchases will be loaded to a data warehouse as the purchases occur.						
Workload type	Updates to inventory data will be loaded to a data warehouse every 1,000 transactions.						
Streaming							

Correct Answer:

Workload Types

Batch
Streaming

Answer Area

Batch
Streaming
Batch

Data for a product catalog will be loaded every 12 hours to a data warehouse.

Data for online purchases will be loaded to a data warehouse as the purchases occur.

Updates to inventory data will be loaded to a data warehouse every 1,000 transactions.

Section:

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 102

DRAG DROP

You have a table named Sales that contains the following data.



SalesDate	SalesAmount	ProductID
4-Apr-20	\$2,000	1
5-Apr-20	\$40	2
5-Apr-20	\$2,300	1
6-Apr-20	\$40	3
6-Apr-20	\$200	4

You need to query the table to return the average sales amount per day. The output must produce the following results.

SalesDate	AVG(SalesAmount)
4-Apr-20	\$2,000
5-Apr-20	\$1,170
6-Apr-20	\$120

How should you complete the query? To answer, drag the appropriate values to the correct targets. Each value 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:

Values	Answer Area
CREATE	Value SalesDate, AVG(SalesAmount)
GROUP BY	FROM Sales
ORDER BY	Value SalesDate
SELECT	ORDER BY SalesDate

Correct Answer:

Values	Answer Area
CREATE	SELECT SalesDate, AVG(SalesAmount)
ORDER BY	FROM Sales
	GROUP BY SalesDate
	ORDER BY SalesDate

Section:

Explanation:

Box 1: SELECT

Box 2: GROUP BY

Example:

When used with a GROUP BY clause, each aggregate function produces a single value covering each group, instead of a single value covering the whole table. The following example produces summary values for each sales territory in the AdventureWorks2012 database. The summary lists the average bonus received by the sales people in each territory, and the sum of year-to-date sales for each territory.

```
SELECT TerritoryID, AVG(Bonus)as 'Average bonus', SUM(SalesYTD) as 'YTD sales'
```

```
FROM Sales.SalesPerson
```

```
GROUP BY TerritoryID;
```

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/functions/avg-transact-sql>

QUESTION 103

DRAG DROP

Match the datastore services to the appropriate descriptions.

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

NOTE: Each correct match is worth one point.

Select and Place:

Services

- Azure Blob storage
- Azure Cosmos DB
- Azure Files
- Azure Table storage

Answer Area

- Service Enables the use of SQL queries against data stored in JSON documents
- Service Enables users to access data by using the Server Message Block (SMB) version 3 protocol

Correct Answer:

Services

- Azure Blob storage
-
-
- Azure Table storage

Answer Area

- Azure Cosmos DB Enables the use of SQL queries against data stored in JSON documents
- Azure Files Enables users to access data by using the Server Message Block (SMB) version 3 protocol

Section:

Explanation:

Box 1: Azure Cosmos DB

In Azure Cosmos DB's SQL (Core) API, items are stored as JSON. The type system and expressions are restricted to deal only with JSON types.

Box 2: Azure Files

Azure Files offers native cloud file sharing services based on the SMB protocol.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql-query-working-with-json>

<https://cloud.netapp.com/blog/azure-smb-server-message-block-in-the-cloud-for-azure-files>

QUESTION 104

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:

Locations

Answer Area

A standalone data analysis tool

Extract:

The CRM system

Load:

The data warehouse

Transform:

Correct Answer:

Locations

Answer Area

Extract:

The CRM system

Load:

The data warehouse

Transform:

A standalone data analysis tool

Section:

Explanation:

Box 1: The CRM system

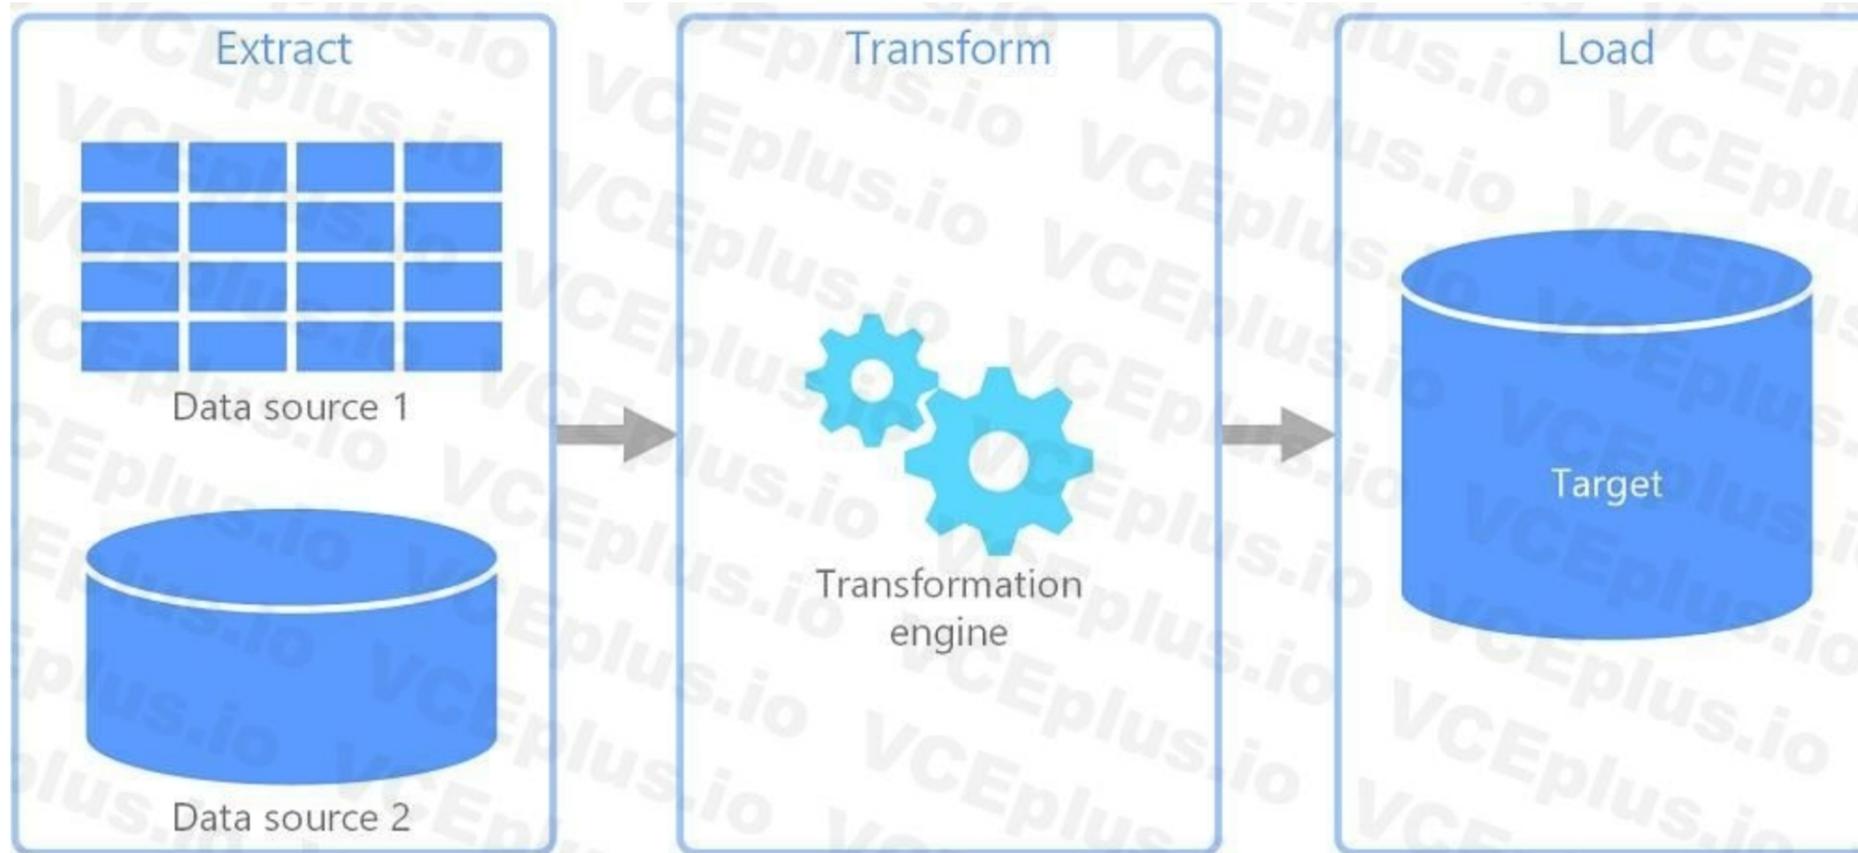
Data is extracted from the CRM system.

Box 2: The data warehouse

Data is loaded to the data warehouse.

Box 3: A standalone data analysis tool

The data transformation that takes place usually involves various operations, such as filtering, sorting, aggregating, joining data, cleaning data, deduplicating, and validating data.



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



QUESTION 105

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:

Analytics Types	Answer Area	
Cognitive	Why did sales increase last month?	<input type="text"/>
Diagnostic	Provide specific guidance on how to allocate current budget to buy different inventory items.	<input type="text"/>
Descriptive	Detect images of the company's logo included in PDF documents	<input type="text"/>
Predictive		
Prescriptive		

Correct Answer:

Analytics Types	Answer Area	
<input type="text"/>	Why did sales increase last month?	Diagnostic
<input type="text"/>	Provide specific guidance on how to allocate current budget to buy different inventory items.	Prescriptive
Descriptive	Detect images of the company's logo included in PDF documents	Cognitive
Predictive		
<input type="text"/>		

Section:

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 patterns, 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.

Box 3: Cognitive

Reference: <https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitiveanalytics/>

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

QUESTION 106

DRAG DROP

Match the data processing languages to the appropriate real-time analytics services.

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

NOTE: Each correct match is worth one point.

Select and Place:

Language	Answer Area
JavaScript	
KQL	
Python	
SQL	

	Azure Data Explorer
	Azure Stream Analytics
	Azure Synapse Analytics

Correct Answer:

Language	Answer Area
JavaScript	

Python	Azure Data Explorer
SQL	Azure Stream Analytics
KQL	Azure Synapse Analytics



Section:

Explanation:

QUESTION 107

DRAG DROP

Match the processes to the appropriate scenarios.

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

NOTE: Each correct match is worth one point.

Select and Place:

Processes	Answer Area
Data ingestion	<input type="text"/> Data is retrieved and processed before being saved to a target system.
Extract, load, and transform (ELT) processing	<input type="text"/> Data is saved to a target system, and then processed.
Extract, transform, and load (ETL) processing	<input type="text"/> Data is extracted from a source system, and then saved to a target system.

Correct Answer:

Processes	Answer Area
<input type="text"/>	Extract, load, and transform (ELT) processing Data is retrieved and processed before being saved to a target system.
<input type="text"/>	Data ingestion Data is saved to a target system, and then processed.
<input type="text"/>	Extract, transform, and load (ETL) processing Data is extracted from a source system, and then saved to a target system.

Section:

Explanation:



QUESTION 108

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:

Term	Answer Area
Data warehouses	<input type="text"/> A solution that copies data from files and online transaction processing (OLTP) databases into an analytical processing system.
Dimension tables	<input type="text"/> Contains numeric values that can be analyzed.
Extract, transform, and load (ETL)	<input type="text"/> A solution that stores data in a relational schema that is optimized for read operations.
Fact tables	
Online analytical processing (OLAP)	

Correct Answer:

Term	Answer Area
	<input type="text" value="Data warehouses"/> A solution that copies data from files and online transaction processing (OLTP) databases into an analytical processing system.
	<input type="text" value="Dimension tables"/> Contains numeric values that can be analyzed.
<input type="text" value="Extract, transform, and load (ETL)"/>	
	<input type="text" value="Fact tables"/> A solution that stores data in a relational schema that is optimized for read operations.
<input type="text" value="Online analytical processing (OLAP)"/>	

Section:
Explanation:

QUESTION 109

DRAG DROP

Match the job roles to the appropriate tasks.

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

NOTE: Each correct match is worth one point.

Select and Place:

Job role	Answer Area
<input type="text" value="Data analyst"/>	<input type="text" value="Database backup and recovery"/>
<input type="text" value="Data engineer"/>	<input type="text" value="Enabling advanced analytic capabilities through reports"/>
<input type="text" value="Database administrator"/>	<input type="text" value="Monitoring data pipelines"/>

Correct Answer:

Job role	Answer Area
<input type="text" value="Database administrator"/>	<input type="text" value="Database backup and recovery"/>
<input type="text" value="Data engineer"/>	<input type="text" value="Enabling advanced analytic capabilities through reports"/>
<input type="text" value="Data analyst"/>	<input type="text" value="Monitoring data pipelines"/>

Section:
Explanation:

QUESTION 110

DRAG DROP

Match the ACID 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 aft.

NOTE: Each correct match is worth one point.

Select and Place:

Terms

- Atomicity
- Consistency
- Isolation
- Durability

Answer Area

- Concurrent transactions cannot interfere with one another and must result in a consistent database state.
- Each transaction is treated as a single unit that succeeds completely or fails completely.
- Transactions can only take the data in a database from one valid state to another.
- When a transaction is committed, it will remain committed.

Correct Answer:

Terms

-
-
-
-

Answer Area

- Isolation Concurrent transactions cannot interfere with one another and must result in a consistent database state.
- Atomicity Each transaction is treated as a single unit that succeeds completely or fails completely.
- Consistency Transactions can only take the data in a database from one valid state to another.
- Durability When a transaction is committed, it will remain committed.

Section:

Explanation: