

Microsoft.DP-900.vMar-2024.by.Isac.111q

Number: DP-900
Passing Score: 800
Time Limit: 120
File Version: 21.0

Website: www.VCEplus.io

Twitter: https://twitter.com/VCE_Plus

Exam Code: DP-900

Exam Name: Microsoft Azure Data Fundamentals



Exam A

QUESTION 1

Which three objects can be added to a Microsoft Power BI dashboard? Each correct answer presents a complete solution. (Choose three.)

NOTE: Each correct selection is worth one point.

- A. a report page
- B. a Microsoft PowerPoint slide
- C. a visualization from a report
- D. a dataflow
- E. a text box

Correct Answer: A, C, E

Section:

Explanation:

Section: Describe an analytics workload on Azure

Reference: <https://docs.microsoft.com/en-us/power-bi/consumer/end-user-dashboards> <https://docs.microsoft.com/en-us/power-bi/create-reports/service-dashboard-add-widget>

QUESTION 2

Which Azure Data Factory component provides the compute environment for activities?

- A. a linked service
- B. an integration runtime
- C. a control flow
- D. a pipeline

Correct Answer: B

Section:

Explanation:

The Integration Runtime (IR) is the compute infrastructure used by Azure Data Factory to provide the following data integration capabilities across different network environments:

Data Flow

Data movement

Activity dispatch

SSIS package execution

Reference:

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

QUESTION 3

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 4

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 5

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 6

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 7

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 8

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 9

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 10

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

- A. Azure Data Lake Storage
- B. Azure Queue storage
- C. Azure Blob storage
- D. Azure Cosmos DB

Correct Answer: A

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 11

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 12

Which clause should you use in a SELECT statement to combine rows in one table with rows in another table?

- A. SET
- B. VALUES
- C. KEY
- D. JOIN

Correct Answer: D

Section:

Explanation:

Reference:

<https://www.tutorialspoint.com/sql/sql-using-joins.htm>

QUESTION 13

Your company needs to ensure that an Azure virtual machine can connect to Azure SQL databases without exposing the databases to the internet. What should you use?

- A. Azure Application Gateway
- B. Azure Traffic Manager
- C. Azure DNS
- D. Azure Private Link

Correct Answer: D

Section:

Explanation:

Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network.

Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary

Reference:

<https://techcommunity.microsoft.com/t5/azure-database-support-blog/azure-sql-db-private-link-private-endpoint-connectivity/ba-p/1235573>

QUESTION 14

What should you use to automatically delete blobs from Azure Blob Storage?

- A. soft delete
- B. archive storage
- C. the change feed
- D. a lifecycle management policy

Correct Answer: D

Section:

QUESTION 15

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 16

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 17

HOTSPOT

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

Hot Area:

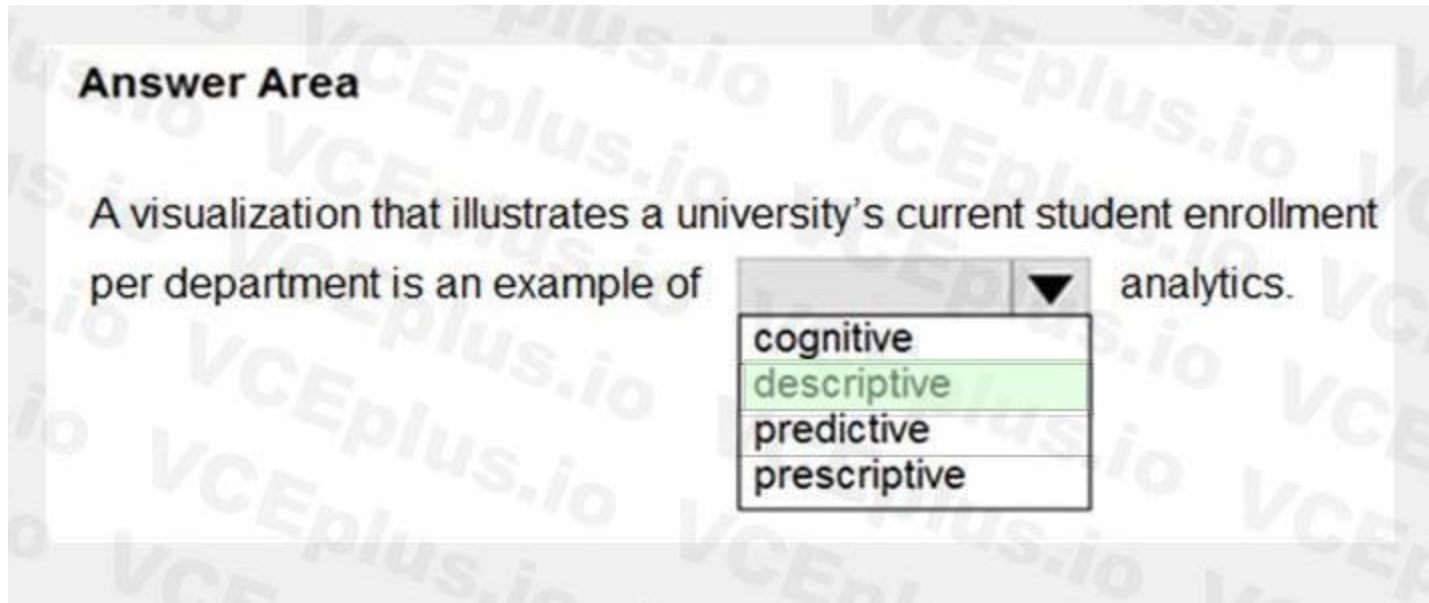
Answer Area

A visualization that illustrates a university's current student enrollment per department is an example of

	▼
cognitive	
descriptive	
predictive	
prescriptive	

 analytics.

Answer Area:



Section:

Explanation:

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

QUESTION 18

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 Table storage within a single Azure Storage account supports multiple concurrent reads in different Azure regions.	<input type="radio"/>	<input type="radio"/>
Azure Table storage within a single Azure Storage account supports multiple concurrent writes in different Azure regions.	<input type="radio"/>	<input type="radio"/>
An Azure Cosmos DB account that uses the Table API supports multiple concurrent reads in different Azure regions.	<input type="radio"/>	<input type="radio"/>
An Azure Cosmos DB account that uses the Table API supports multiple concurrent writes in different Azure regions.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
Azure Table storage within a single Azure Storage account supports multiple concurrent reads in different Azure regions.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Table storage within a single Azure Storage account supports multiple concurrent writes in different Azure regions.	<input type="radio"/>	<input checked="" type="radio"/>
An Azure Cosmos DB account that uses the Table API supports multiple concurrent reads in different Azure regions.	<input checked="" type="radio"/>	<input type="radio"/>
An Azure Cosmos DB account that uses the Table API supports multiple concurrent writes in different Azure regions.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

Box 1: Yes 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).

Box 2: No

Box 3: Yes

Box 4: Yes

Azure Cosmos DB supports multi-region writes.

Reference:

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

<https://manojchoudhari.wordpress.com/2019/12/16/azure-cosmos-db-enable-multi-region-writes>

QUESTION 19

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 access control lists that have object-level permissions,

	▼
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 access control lists that have object-level permissions,

	▼
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:

A key mechanism that allows Azure Data Lake Storage Gen2 to provide file system performance at object storage scale and prices is the addition of a hierarchical namespace. This allows the collection of objects/files within an account to be organized into a hierarchy of directories and nested subdirectories in the same way that the file system on your computer is organized. With a hierarchical namespace enabled, a storage account becomes

capable of providing the scalability and cost-effectiveness of object storage, with file system semantics that are familiar to analytics engines and frameworks.
Reference:
<https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-namespaces>

QUESTION 20
You have a quality assurance application that reads data from a data warehouse.
Which type of processing does the application use?

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

Correct Answer: A
Section:
Explanation:
Section: Describe an analytics workload on Azure
Reference: <https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/online-transaction-processing>

QUESTION 21
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
You can copy a dashboard between Microsoft Power BI workspaces.	<input type="radio"/>	<input type="radio"/>
A Microsoft Power BI dashboard can only display visualizations from a single dataset.	<input type="radio"/>	<input type="radio"/>
A Microsoft Power BY dashboard can display visualizations from a Microsoft Excel workbook.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Statements	Yes	No
You can copy a dashboard between Microsoft Power BI workspaces.	<input type="radio"/>	<input checked="" type="radio"/>
A Microsoft Power BI dashboard can only display visualizations from a single dataset.	<input type="radio"/>	<input checked="" type="radio"/>
A Microsoft Power BY dashboard can display visualizations from a Microsoft Excel workbook.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

Box 1: No

You can duplicate a dashboard. The duplicate ends up in the same Power BI workspace.

There is no current functionality that allows you to move reports from one workspace to another.

Box 2: No

Box 3: Yes

Reference:

<https://docs.microsoft.com/en-us/power-bi/connect-data/service-datasets-across-workspaces>

<https://docs.microsoft.com/en-us/power-bi/consumer/end-user-dashboards><https://powerbi.microsoft.com/en-us/excel-and-power-bi/>

QUESTION 22

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

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

Answer Area:

Answer Area

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

A data analyst

A data engineer

A data scientist

A database administrator

Section:

Explanation:

QUESTION 23

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:

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.

Section:

Explanation:

QUESTION 24

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A relational database must be used when

- strong consistency guarantees are required.
- 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

- strong consistency guarantees are required.
- 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 25

You need to store event log data that is semi-structured and received as the logs occur. What should you use?

- A. Azure Table storage
- B. Azure Queue storage
- C. Azure Files

Correct Answer: C

Section:

Explanation:

QUESTION 26

What is a characteristic of a non-relational database?

- A. a fixed schema
- B. self-describing entities
- C. full support for Transact-SQL

Correct Answer: B

Section:

QUESTION 27

Select the answer that correctly completes the sentence.

Answer Area

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

A. See the answer in explanation below

Correct Answer: A

Section:

Explanation:

Answer Area

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

QUESTION 28

Select the answer that correctly completes the sentence.

Answer Area

When provisioning an Azure Cosmos DB , you need to specify which type of API you will use.

A. See the answer in explanation below

Correct Answer: A

Section:

Explanation:

Answer Area

When provisioning an Azure Cosmos DB , you need to specify which type of API you will use.

QUESTION 29

When you create an Azure SQL database, which account can always connect to the database?

- A. the Azure Active Directory (Azure AD) account that created the database
- B. the server admin login account of the logical server
- C. the Azure Active Directory (Azure AD) administrator account

D. the sa account

Correct Answer: B

Section:

Explanation:

Section: Describe how to work with relational data on Azure

When you first deploy Azure SQL, you specify an admin login and an associated password for that login. This administrative account is called Server admin.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-quickstart>

QUESTION 30

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 31

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 32

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 33

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 34

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:

www.VCEplus.io

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 35

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 36

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 37

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 38

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 39

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 40

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 41

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 42

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 43

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)

www.VCEplus.io

- 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 44

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 45

Select the answer the correctly completes the sentence.

Answer Area

is responsible for managing the security of the data in a database.

- A. See the answer in explanation below

Correct Answer: A

Section:

Explanation:

Answer Area

is responsible for managing the security of the data in a database.

QUESTION 46

Select the answer that correctly completes the sentence.

Answer Area

A block of code that runs in a database is called

A. See the answer in explanation below

Correct Answer: A

Section:

Explanation:

Answer Area

A block of code that runs in a database is called

www.VCEplus.io

QUESTION 47

Select the answer that correctly completes the sentence.

Answer Area

The clause can be used in Data Manipulation Language (DML) statements to specify the criteria that rows must match.

A. See the answer in explanation below

Correct Answer: A

Section:

Explanation:

Answer Area

The clause can be used in Data Manipulation Language (DML) statements to specify the criteria that rows must match.

QUESTION 48

You have data saved in the following format.

```
{
  "firstName": "John",
  "lastName": "Doe",
  "age": "23",
  "hobbies": [
    { "type": "Sports", "value": "Golf" },
    { "type": "Leisure", "value": "Reading" },
    { "type": "Leisure", "value": "Guitar" }
  ]
}
```

Which format was used?

- A. XML
- B. yaml
- C. JSON
- D. HTML

Correct Answer: A

Section:

QUESTION 49

What is a primary characteristic of a relational database?

- A. data is queried and manipulated by using a variant of the SQL language
- B. a flexible data structure
- C. a lack of dependencies between tables
- D. a large amount of duplicate data

Correct Answer: A

Section:

QUESTION 50

Which database transaction property ensures that individual transactions are executed only once and either succeed in their entirety or roll back?

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

Correct Answer: C

Section:

QUESTION 51

Select the answer that correctly completes the sentence.

Answer Area

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

- A. See the answer in explanation below

Correct Answer: A

Section:

Explanation:

Answer Area

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

QUESTION 52

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

The format is used to store delimited data, separates data fields by using commas, and terminates rows by using a carriage return

- CSV
- JSON
- TSV
- XML

Answer Area:

Answer Area

The CSV format is used to store delimited data, separates data fields by using commas, and terminates rows by using a carriage return

CSV
JSON
TSV
XML

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 53

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

The JSON file format uses a hierarchical document schema to define data entities that have multiple attributes.

CSV
JSON
ORC
TSV

Answer Area:

Answer Area

The JSON file format uses a hierarchical document schema to define data entities that have multiple attributes.

CSV
JSON
ORC
TSV

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 54

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

Database administrators apply data cleansing routines and turn data into useful information.	Yes <input type="radio"/>	No <input type="radio"/>
Data engineers manage databases, store backup copies of data, and restore data in the event of a failure.	<input type="radio"/>	<input type="radio"/>
Data analysts create data visuals and enable companies to make data-driven decisions.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements

Database administrators apply data cleansing routines and turn data into useful information.	Yes <input type="radio"/>	No <input checked="" type="radio"/>
Data engineers manage databases, store backup copies of data, and restore data in the event of a failure.	<input type="radio"/>	<input checked="" type="radio"/>
Data analysts create data visuals and enable companies to make data-driven decisions.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

QUESTION 55

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A database administrator is responsible for managing the security of the data in a database.

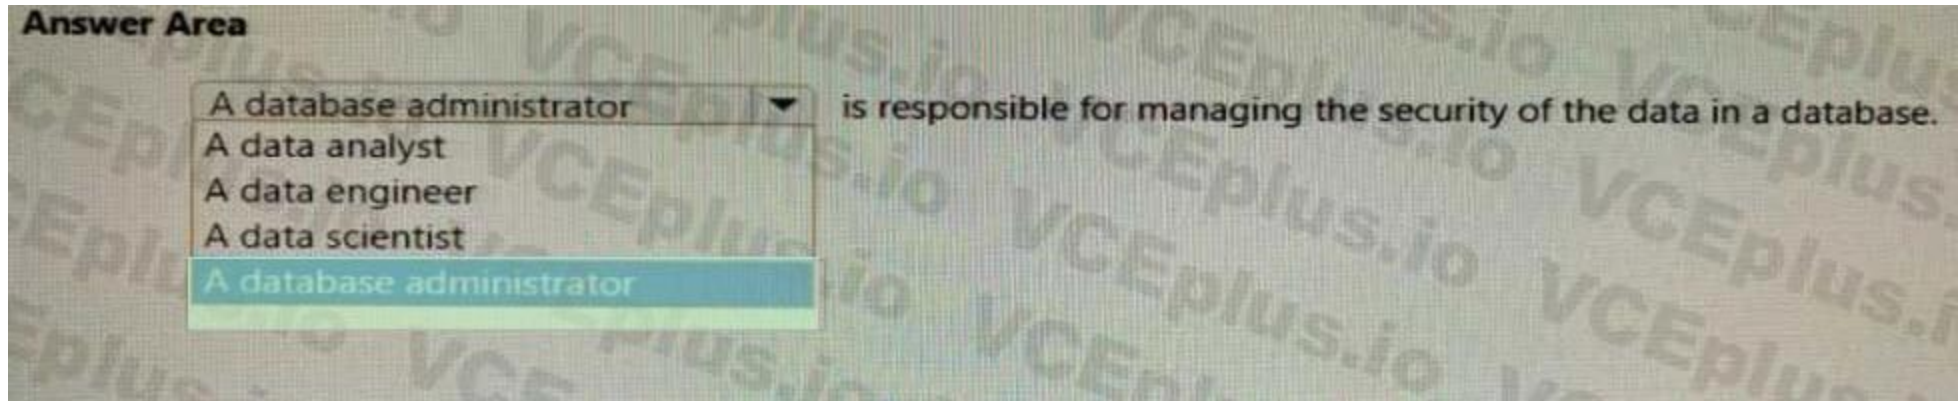
A data analyst

A data engineer

A data scientist

A database administrator

Answer Area:



Section:

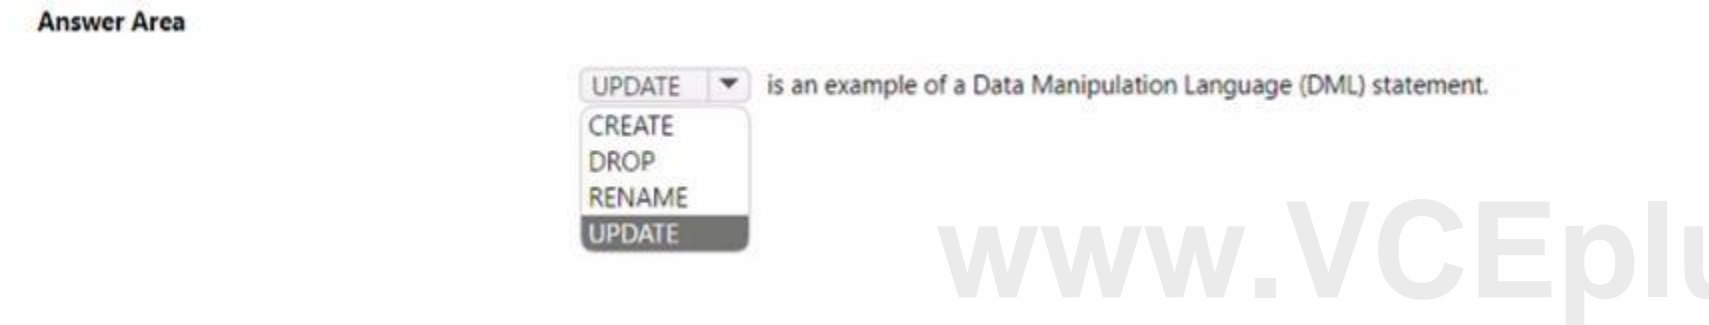
Explanation:

QUESTION 56

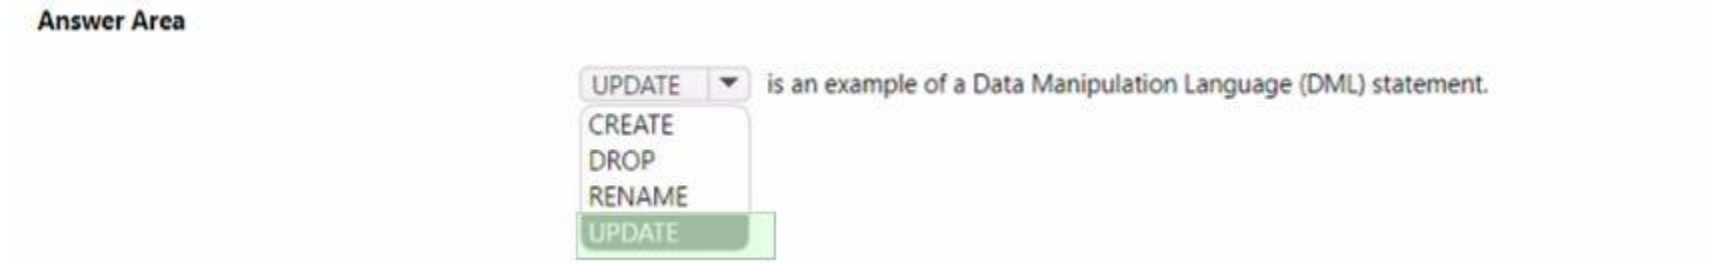
HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



Answer Area:

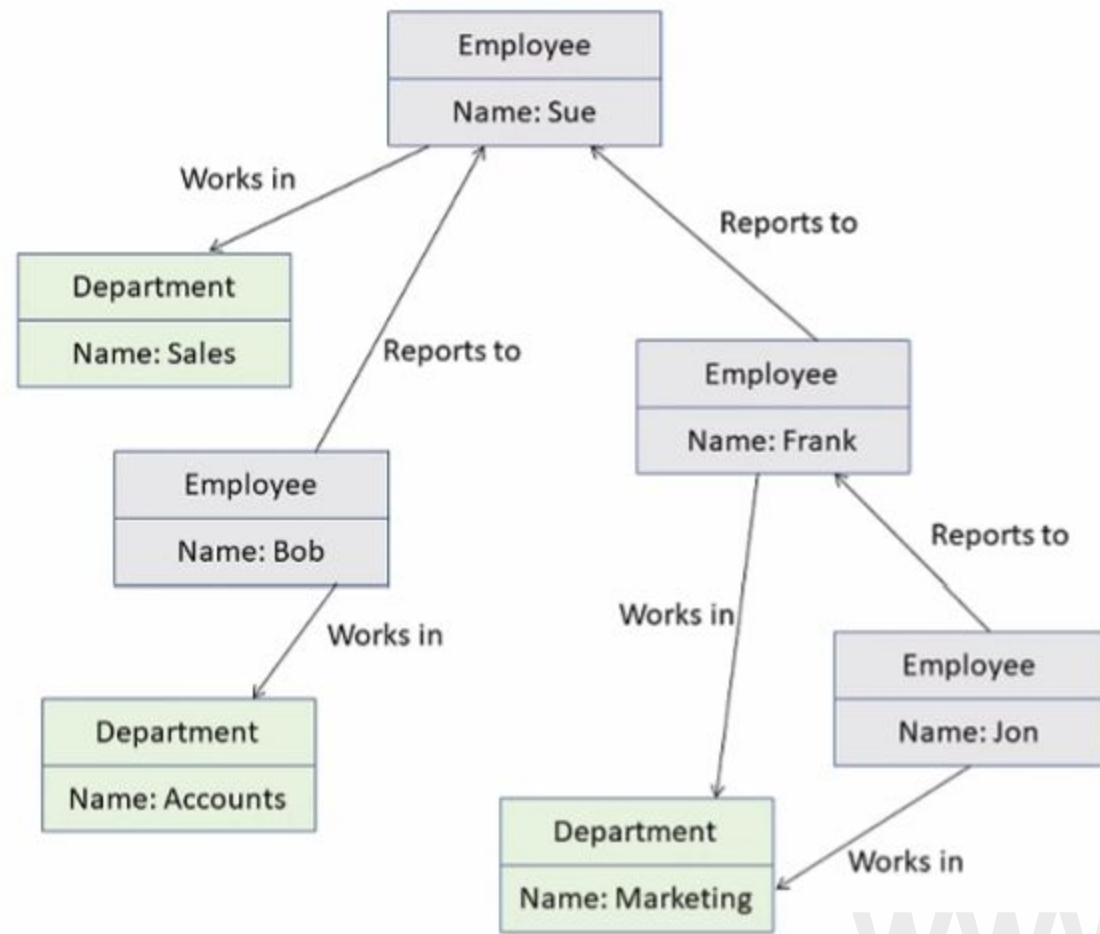


Section:

Explanation:

QUESTION 57

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



Which type of data store is this?

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

Correct Answer: B

Section:

QUESTION 58

In an analytical data model, which type of table contains entities that are used to aggregate numeric values, where each entity is represented by a row with a unique key value?

- A. bridge
- B. dimension
- C. fact

Correct Answer: C

Section:

QUESTION 59

Which Azure Cosmos DB API should you use for a graph database?

- A. Azure Cosmos DB for Apache Gremlin
- B. Azure Cosmos DB for NoSQL
- C. Azure Cosmos DB for Apache Cassandra
- D. Azure Cosmos DB for Table

Correct Answer: A
Section:

QUESTION 60
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 a transactional database that guarantees atomicity, all statements in a transaction must complete successfully or the transaction must roll back.	<input type="radio"/>	<input type="radio"/>
In a transactional database that guarantees consistency, all data changes will be logged and accounted for.	<input type="radio"/>	<input type="radio"/>
In a transactional database that guarantees isolation, processes can see only data in a consistent state.	<input type="radio"/>	<input type="radio"/>

Answer Area:
Answer Area

Statements	Yes	No
In a transactional database that guarantees atomicity, all statements in a transaction must complete successfully or the transaction must roll back.	<input type="radio"/>	<input checked="" type="radio"/>
In a transactional database that guarantees consistency, all data changes will be logged and accounted for.	<input type="radio"/>	<input checked="" type="radio"/>
In a transactional database that guarantees isolation, processes can see only data in a consistent state.	<input checked="" type="radio"/>	<input type="radio"/>

Section:
Explanation:

QUESTION 61
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 62

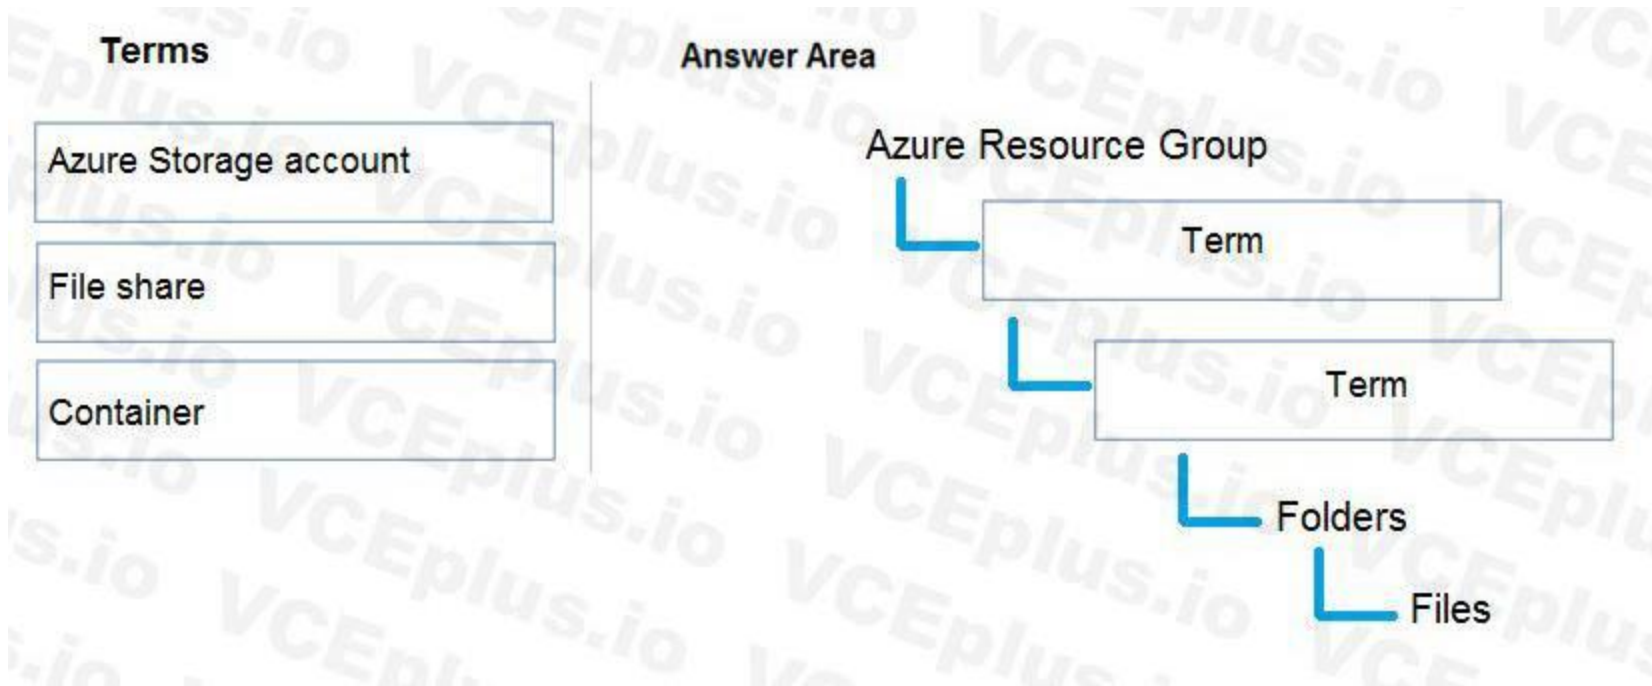
DRAG DROP

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

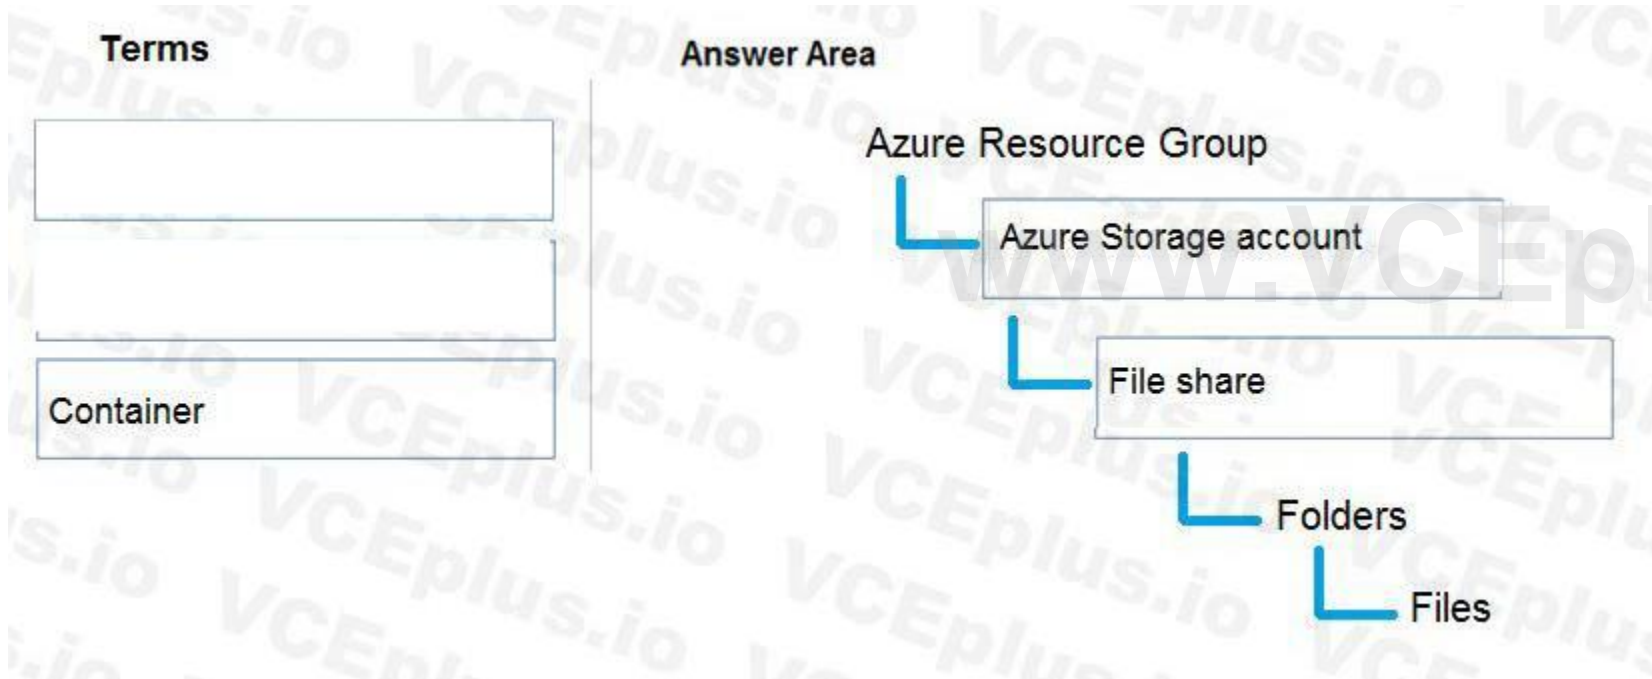
To answer, drag the appropriate term from the column on the left to its level on the right. Each term may be used once, more than once, or not at all.

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 63

In a fully denormalized database, how is data read and written for a single entity?

- A. Data is read from a single table and written to a single table.
- B. Data is read from multiple tables and written to a single table.

- C. Data is read from a single table and written to multiple tables.
- D. Data is read from multiple tables and written to multiple tables.

Correct Answer: B

Section:

QUESTION 64

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

Correct Answer: D

Section:

QUESTION 65

You have a banking application that transfers money in to and out of accounts.

Of which type of solution is this is an example?

- A. extract transform, and load (ETL)
- B. a data warehouse
- C. online analytical processing (OLAP)
- D. online transaction processing (OLTP)

Correct Answer: D

Section:

QUESTION 66

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

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

Correct Answer: D

Section:

QUESTION 67

Select the answer that correctly completes the sentence.

www.VCEplus.io

Answer Area

A JSON document is an example of

A. See the answer in explanation below

Correct Answer: A
Section:
Explanation:

Answer Area

A JSON document is an example of

QUESTION 68
HOTSPOT
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Hot Area:

Statements	Yes	No
A pipeline is a representation of a data structure within Azure Data Factory	<input type="radio"/>	<input type="radio"/>
Azure Data Factory pipelines can execute other pipelines	<input type="radio"/>	<input type="radio"/>
A processing step within an Azure Data Factory pipeline is an activity	<input type="radio"/>	<input type="radio"/>

Answer Area:

Statements	Yes	No
A pipeline is a representation of a data structure within Azure Data Factory	<input type="radio"/>	<input checked="" type="radio"/>
Azure Data Factory pipelines can execute other pipelines	<input checked="" type="radio"/>	<input type="radio"/>
A processing step within an Azure Data Factory pipeline is an activity	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

Box 1: No

A pipeline is a logical grouping of activities that together perform a task.

Box 2: Yes

You can construct pipeline hierarchies with data factory.

Box 3: Yes

A pipeline is a logical grouping of activities that together perform a task.

Reference:

<https://mrpaulandrew.com/2019/09/25/azure-data-factory-pipeline-hierarchies-generation-control/>

QUESTION 69

What should you use to build a Microsoft Power BI paginated report?

- A. Charticator
- B. Power BI Desktop
- C. the Power BI service
- D. Power BI Report Builder

Correct Answer: D

Section:

Explanation:

Section: Describe an analytics workload on Azure

Power BI Report Builder is the standalone tool for authoring paginated reports for the Power BI service.

Reference:

<https://docs.microsoft.com/en-us/power-bi/paginated-reports/paginated-reports-report-builder-power-bi>

QUESTION 70

You have a SQL query that combines customer data and order data. The query includes calculated columns.

You need to persist the SQL query so that other users can use the query.

What should you create?

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

Correct Answer: B

Section:

Explanation:

Section: Describe core data concepts

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 71

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

Correct Answer: B, D

Section:

Explanation:

Section: Describe core data concepts

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 72

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

Correct Answer: C

Section:

Explanation:

Section: Describe how to work with relational data on Azure

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

www.VCEplus.io

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

QUESTION 73

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

Correct Answer: A, C

Section:

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:

Statements	Yes	No
If you have a platform as a service (PaaS) database in Azure, you are responsible for applying operating system updates.	<input type="radio"/>	<input type="radio"/>
If you have a platform as a service (PaaS) database in Azure, backups are performed automatically.	<input type="radio"/>	<input type="radio"/>
If you have a platform as a service (PaaS) database in Azure, you are responsible for upgrading the database engine.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Statements	Yes	No
If you have a platform as a service (PaaS) database in Azure, you are responsible for applying operating system updates.	<input type="radio"/>	<input checked="" type="radio"/>
If you have a platform as a service (PaaS) database in Azure, backups are performed automatically.	<input checked="" type="radio"/>	<input type="radio"/>
If you have a platform as a service (PaaS) database in Azure, you are responsible for upgrading the database engine.	<input type="radio"/>	<input checked="" type="radio"/>

Section:

Explanation:

Box 1: No

Microsoft handles all patching and updating of the SQL and operating system code. You don't have to manage the underlying infrastructure.

Box 2: Yes

SQL Database is a fully managed service that has built-in high availability, backups, and other common maintenance operations.

Box 3: No

Reference:

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

QUESTION 75

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 Databricks can consume data from Azure SQL Database	<input type="radio"/>	<input type="radio"/>
Azure Databricks can consume data from Azure Event Hubs	<input type="radio"/>	<input type="radio"/>
Azure Databricks can consume data from Azure Cosmos DB	<input type="radio"/>	<input type="radio"/>

Answer Area:

Statements	Yes	No
Azure Databricks can consume data from Azure SQL Database	<input checked="" type="radio"/>	<input type="radio"/>
Azure Databricks can consume data from Azure Event Hubs	<input checked="" type="radio"/>	<input type="radio"/>
Azure Databricks can consume data from Azure Cosmos DB	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

Box 1: Yes

Azure Databricks can consume data from SQL Databases using JDBC and from SQL Databases using the Apache Spark connector.

The Apache Spark connector for Azure SQL Database and SQL Server enables these databases to act as input data sources and output data sinks for Apache Spark jobs.

Box 2: Yes

You can stream data into Azure Databricks using Event Hubs.

Box 3: Yes You can run Spark jobs with data stored in Azure Cosmos DB using the Cosmos DB Spark connector. Cosmos can be used for batch and stream processing, and as a serving layer for low latency access. You can use the connector with Azure Databricks or Azure HDInsight, which provide managed Spark clusters on Azure.

Reference:

<https://docs.microsoft.com/en-us/azure/databricks/data/data-sources/sql-databases-azure>

<https://docs.microsoft.com/en-us/azure/databricks/scenarios/databricks-stream-from-eventhubs>

QUESTION 76

HOTSPOT

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

Hot Area:

Answer Area

A Microsoft Power BI

▼
dashboard
interactive report
paginated report
subscription

 enables users to create highly formatted,

fixed-layout documents optimized for printing and archiving.

Answer Area:

Answer Area

A Microsoft Power BI

▼
dashboard
interactive report
paginated report
subscription

 enables users to create highly formatted,

fixed-layout documents optimized for printing and archiving.

Section:

Explanation:

Paginated Reports in Power BI now allows users to generate these fixed-layout documents optimized for printing and archiving, such as PDF and Word files. These document-style reports with visualizations that provide additional control, like which tables expand horizontally and vertically to display all their data and continue from page to page as needed.

Reference:
<https://powerbi.microsoft.com/en-us/blog/announcing-paginated-reports-in-power-bi-general-availability/>

QUESTION 77
HOTSPOT
Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

Database administrators

Data analysts

Data engineers

Data scientists

Database administrators

are responsible for the availability, performance, and optimization of databases.

Answer Area:

Answer Area

Database administrators

Data analysts

Data engineers

Data scientists

Database administrators

are responsible for the availability, performance, and optimization of databases.

Section:
Explanation:

QUESTION 78
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

Dataset

Linked service

Mapping data flow

Pipeline

Answer Area

Component

A representation of data structures within data stores

Component

The information used to connect to external resources

Component

A logical grouping of activities that performs a unit of work and can be scheduled

Correct Answer:

Components

Mapping data flow

Answer Area

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

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 79

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

Correct Answer:

Workload Types	Answer Area
Batch	Batch Data for a product catalog will be loaded every 12 hours to a data warehouse.
Streaming	Streaming Data for online purchases will be loaded to a data warehouse as the purchases occur.
	Batch 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 80

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

CREATE

ORDER BY

GROUP BY

SELECT

ANSWER AREA

ValueSalesDate, AVG(SalesAmount)

FROM Sales

ValueSalesDate

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>

www.VCEplus.io

QUESTION 81

HOTSPOT

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

Hot Area:

Answer Area
When provisioning an Azure Cosmos DB <div> <div></div> <div> account container database item </div> </div> , you need to specify which type of API you will use.

Answer Area:

Answer Area

When provisioning an Azure Cosmos DB

	▼
account	
container	
database	
item	

, you need to specify which type of API you will use.

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/create-cosmosdb-resources-portal>

QUESTION 82

HOTSPOT

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

www.VCEplus.io

Hot Area:

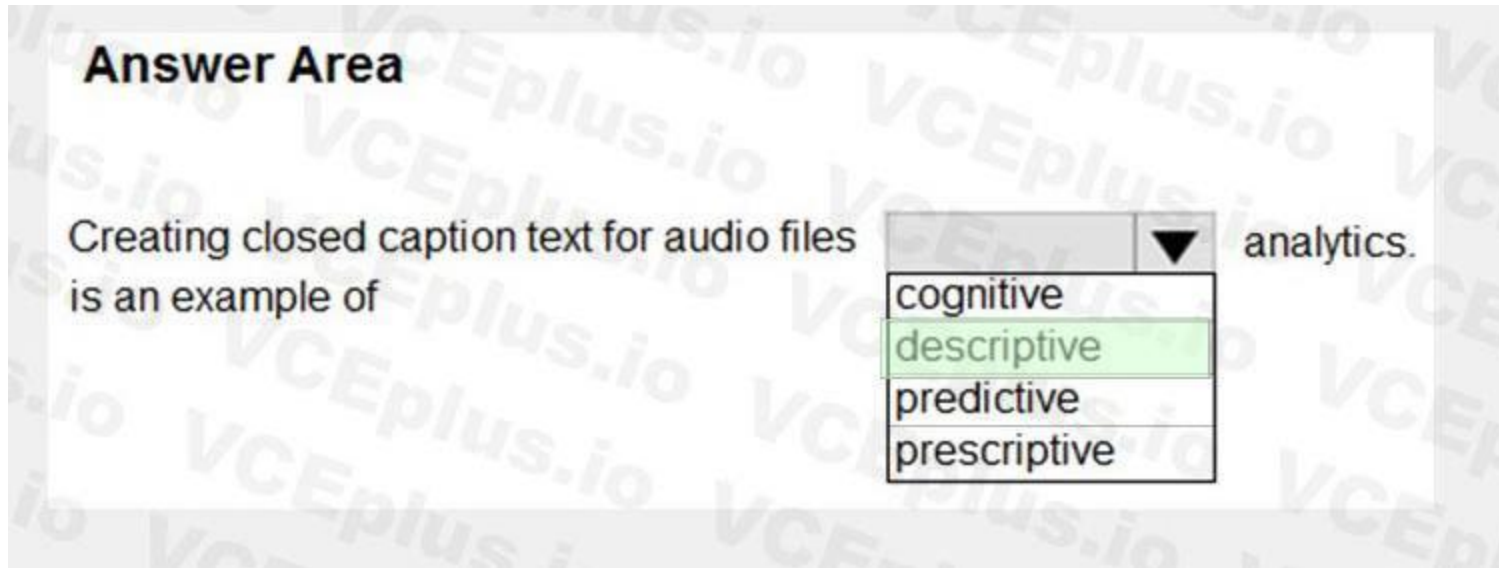
Answer Area

Creating closed caption text for audio files is an example of

	▼
cognitive	
descriptive	
predictive	
prescriptive	

 analytics.

Answer Area:



Section:

Explanation:

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



Note: Azure Media Indexer enables you to make content of your media files searchable and to generate a full-text transcript for closed captioning and keywords. You can process one media file or multiple media files in a batch.

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/>

<https://docs.microsoft.com/en-us/azure/media-services/previous/media-services-index-content>

QUESTION 83

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

Answer 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"/>

Section:

Explanation:

QUESTION 84

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.

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.

The Microsoft Power BI service

The Microsoft Power BI phone app

Section:

Explanation:

www.VCEplus.io

QUESTION 85

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

can be used as part of a lift-and-shift solution for migrating
 changes.

Answer Area:

Answer Area

can be used as part of a lift-and-shift solution for migrating
 changes.

Section:

Explanation:

Answer Area

can be used as part of a lift-and-shift solution for migrating
 changes.

QUESTION 86

Select the answer that correctly completes the sentence.

Answer Area

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

A. See the answer in explanation below

Correct Answer: A

Section:

Explanation:

Answer Area

The massively parallel processing (MPP) engine
of Azure Synapse Analytics distributes processing across compute nodes.

QUESTION 87

Select the answer that correctly completes the sentence.

Answer Area

can be used as part of a lift-and-shift solution for migrating
an on-premises application to the cloud with minimal changes.

A. See the answer in explanation below

Correct Answer: A

Section:

Explanation:

Answer Area

can be used as part of a lift-and-shift solution for migrating
an on-premises application to the cloud with minimal changes.

QUESTION 88

What can be used with native notebook support to query and visualize data by using a web-based interface?

- A. Azure Databricks
- B. pgAdmin
- C. Microsoft Power BI

Correct Answer: A

Section:

Explanation:

QUESTION 89

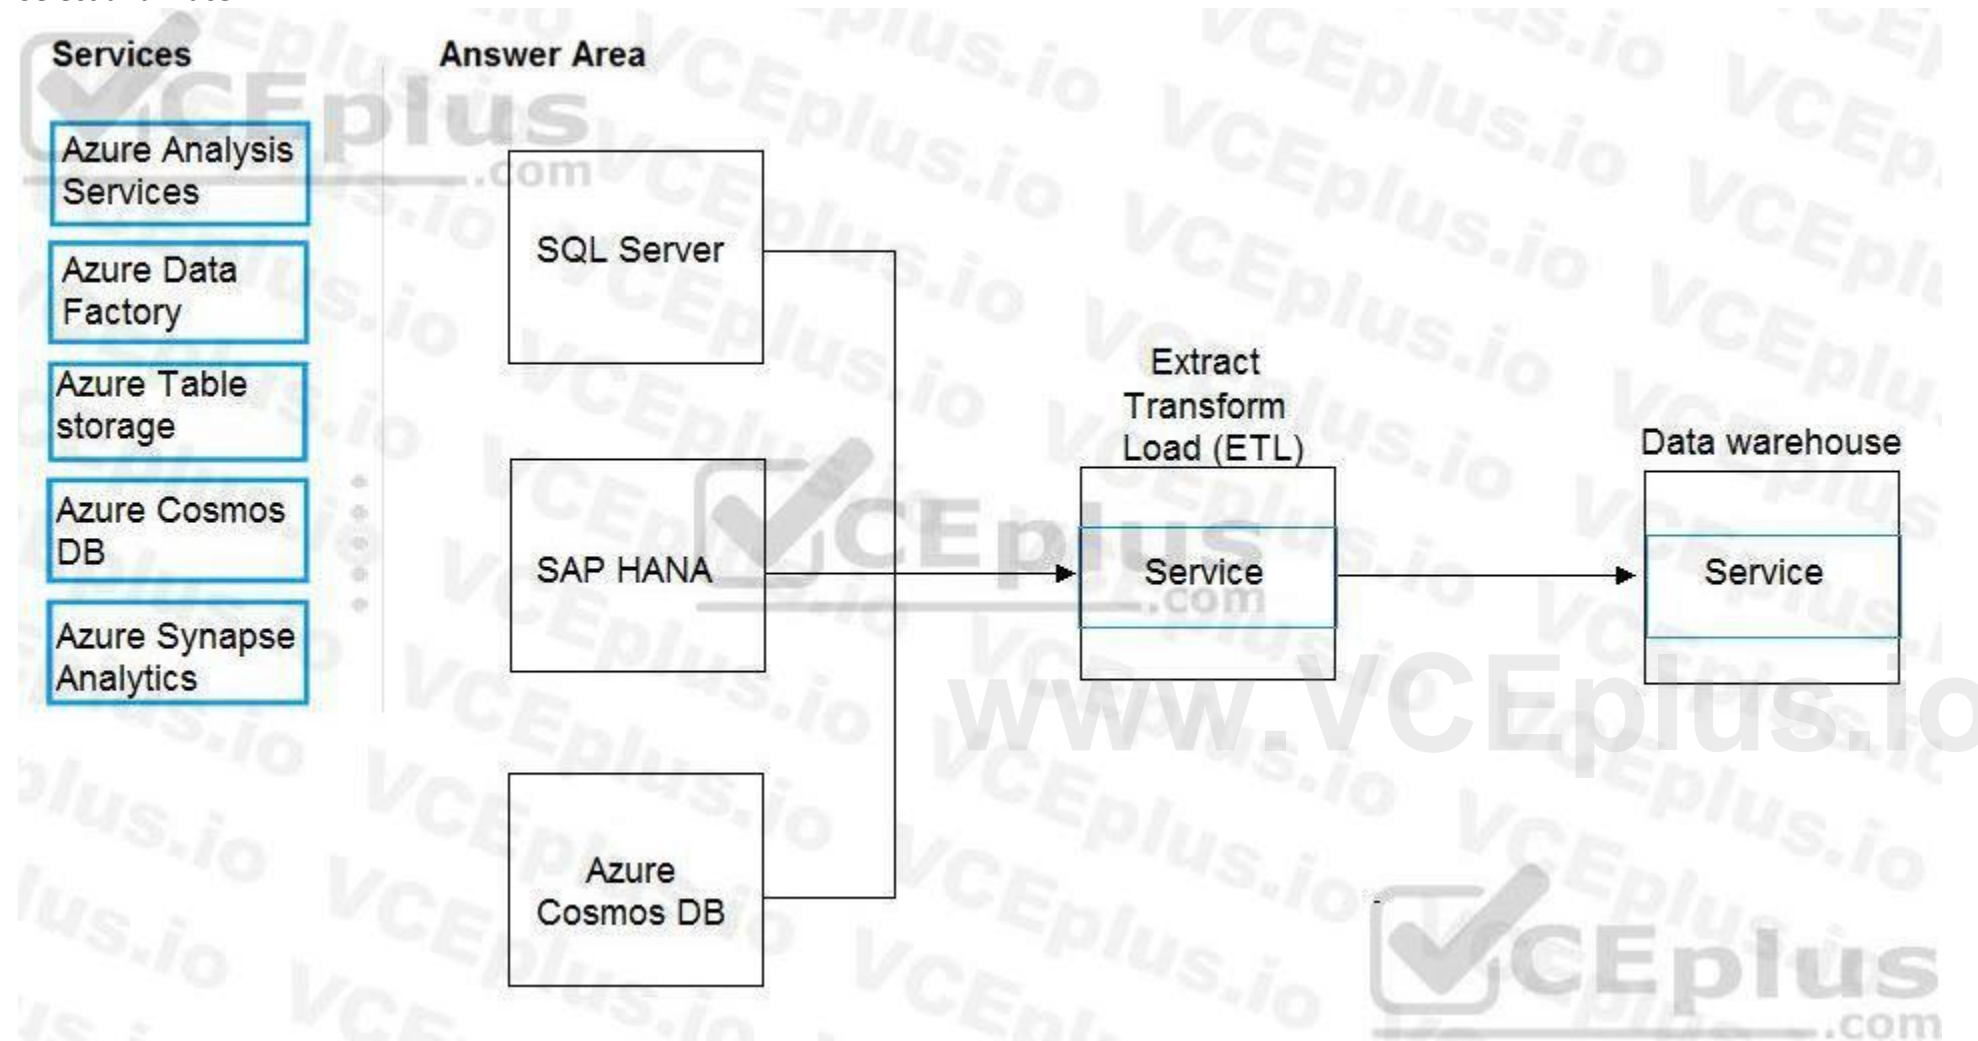
DRAG DROP

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

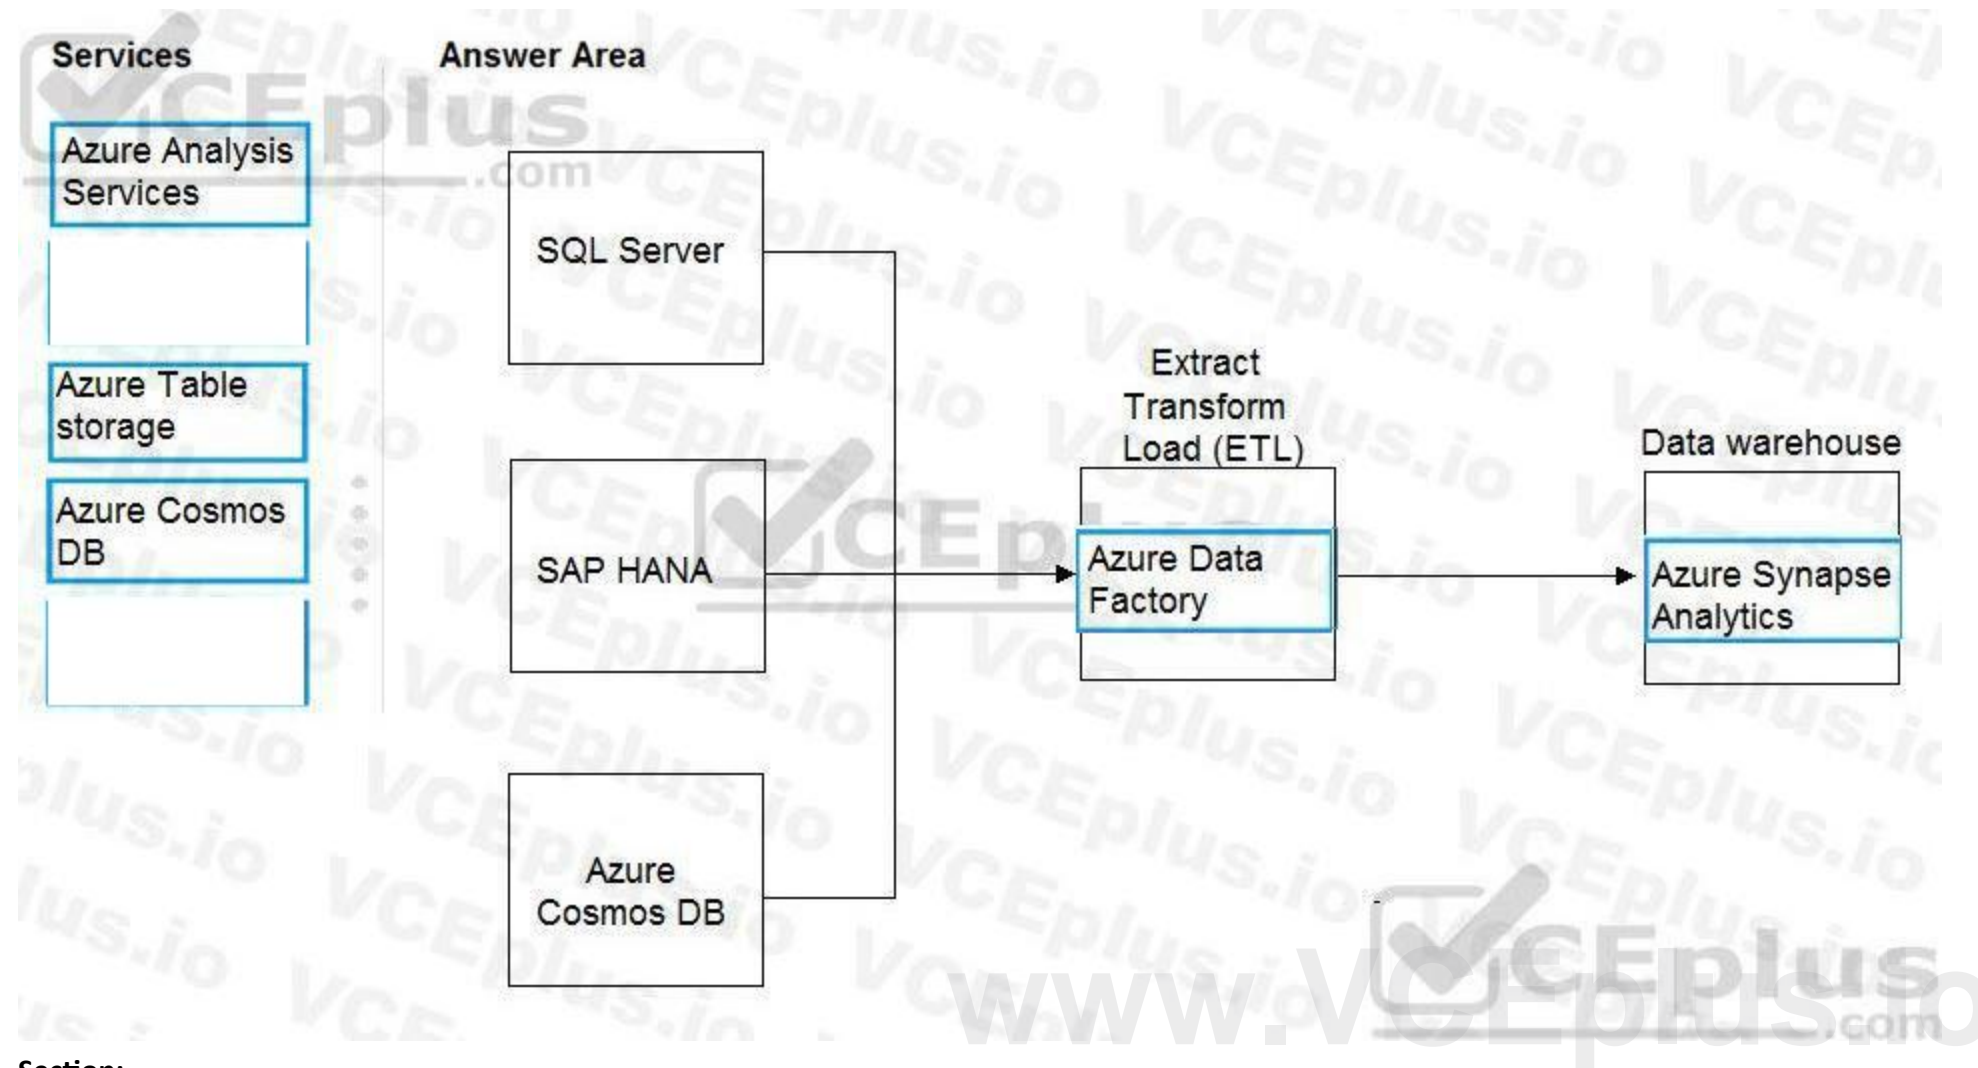
To answer, drag the appropriate service from the column on the left to its location on the right. Each service may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:



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 90

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A data analyst is responsible for identifying which business rules must be applied to the data of a company.

A data analyst

A data engineer

A data scientist

Answer Area:

Answer Area

A data analyst is responsible for identifying which business rules must be applied to the data of a company.

A data analyst

A data engineer

A data scientist

Section:

Explanation:

QUESTION 91

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 92

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 Synapse Analytics scales storage and compute independently	<input type="radio"/>	<input type="radio"/>
Azure Synapse Analytics can be paused to reduce compute costs	<input type="radio"/>	<input type="radio"/>
An Azure Synapse Analytics data warehouse has a fixed storage capacity	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area	Statements	Yes	No
	Azure Synapse Analytics scales storage and compute independently	<input checked="" type="radio"/>	<input type="radio"/>
	Azure Synapse Analytics can be paused to reduce compute costs	<input checked="" type="radio"/>	<input type="radio"/>
	An Azure Synapse Analytics data warehouse has a fixed storage capacity	<input type="radio"/>	<input checked="" type="radio"/>

Section:

Explanation:

Box 1: Yes

Compute is separate from storage, which enables you to scale compute independently of the data in your system.

Box 2: Yes

You can use the Azure portal to pause and resume the dedicated SQL pool compute resources.

Pausing the data warehouse pauses compute. If your data warehouse was paused for the entire hour, you will not be charged compute during that hour.

Box 3: No

Storage is sold in 1 TB allocations. If you grow beyond 1 TB of storage, your storage account will automatically grow to 2 TBs.

Reference:

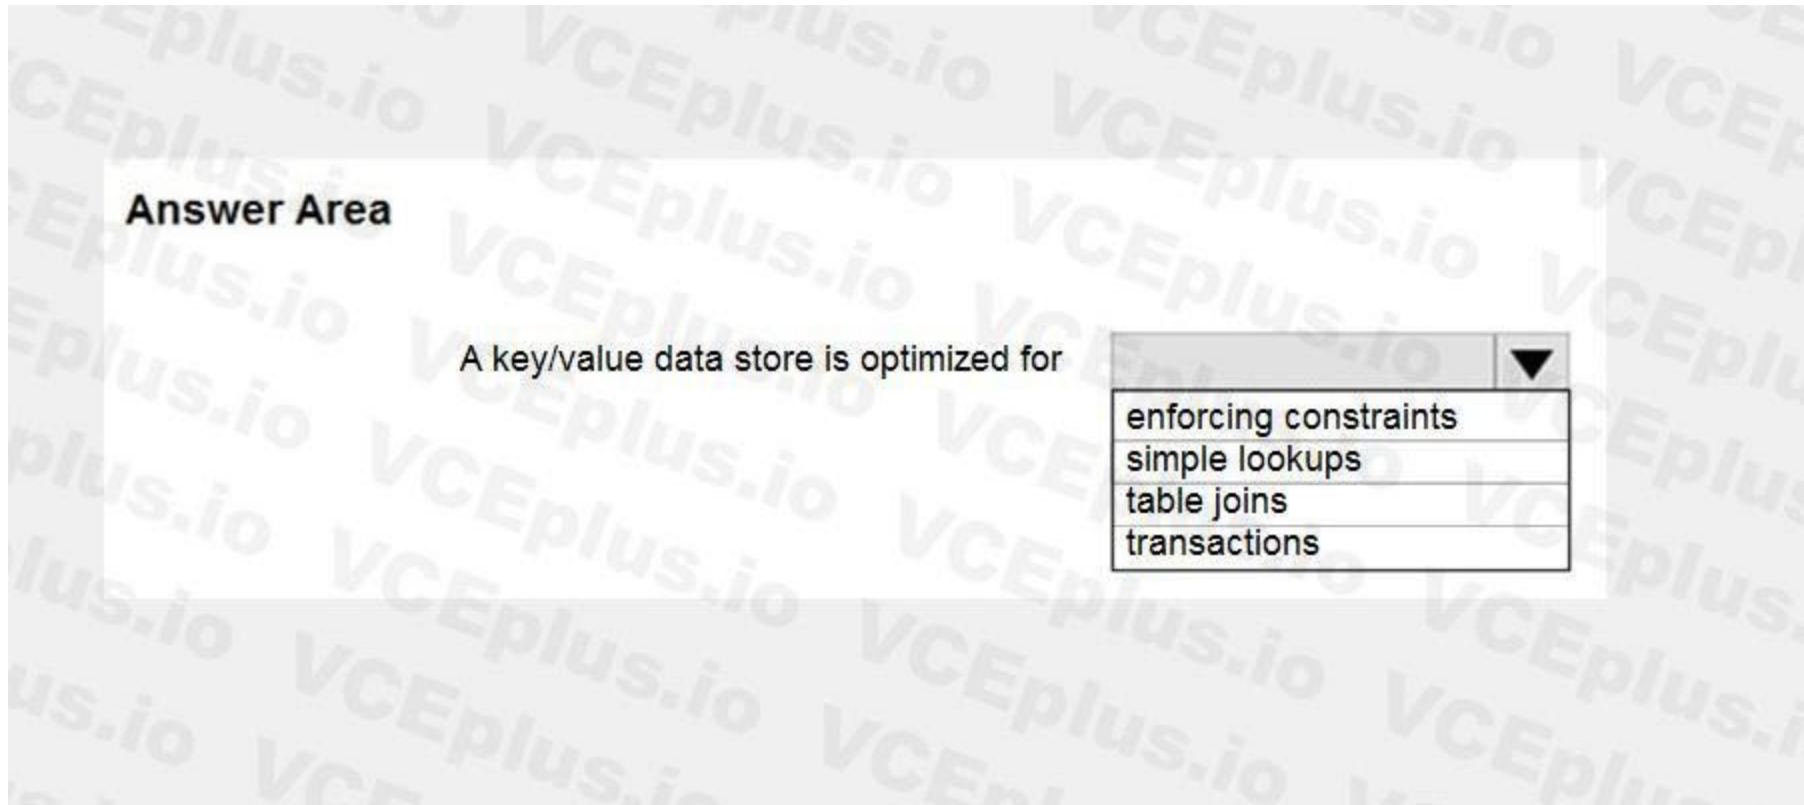
<https://azure.microsoft.com/en-us/pricing/details/synapse-analytics/>

QUESTION 93

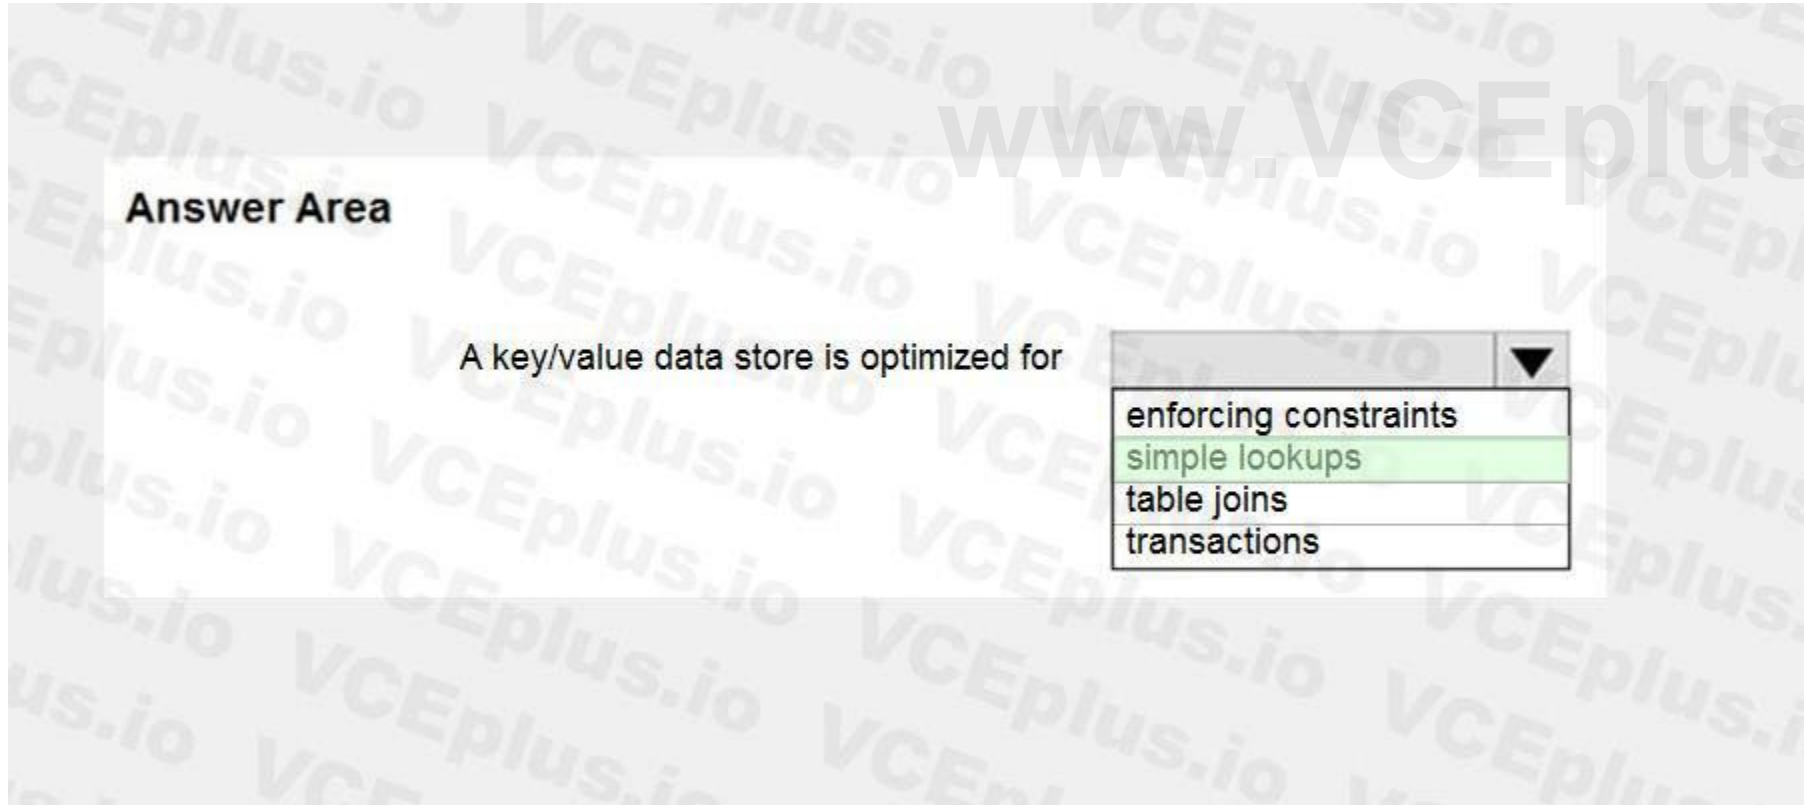
HOTSPOT

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

Hot Area:



Answer Area:



Section:

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 94

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 Data Studio can be used to query an Azure SQL database from a device that runs macOS.	<input type="radio"/>	<input type="radio"/>
Microsoft SQL Server Management Studio (SSMS) enables users to create and use SQL notebooks.	<input type="radio"/>	<input type="radio"/>
Azure Data Studio can be used to restore a database.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Statements	Yes	No
Azure Data Studio can be used to query an Azure SQL database from a device that runs macOS.	<input checked="" type="radio"/>	<input type="radio"/>
Microsoft SQL Server Management Studio (SSMS) enables users to create and use SQL notebooks.	<input type="radio"/>	<input checked="" type="radio"/>
Azure Data Studio can be used to restore a database.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

Box 1: No

Microsoft handles all patching and updating of the SQL and operating system code. You don't have to manage the underlying infrastructure.

Box 2: Yes

SQL Database is a fully managed service that has built-in high availability, backups, and other common maintenance operations.

Box 3: No

Reference:

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

QUESTION 95

You need to use Transact-SQL to query files in Azure Data Lake Storage from an Azure Synapse Analytics data warehouse.

What should you use to query the files?

- A. Azure Functions
- B. Microsoft SQL Server Integration Services (SSIS)
- C. PolyBase
- D. Azure Data Factory

Correct Answer: C

Section:

Explanation:

Section: Describe an analytics workload on Azure

Reference: <https://docs.databricks.com/data/data-sources/azure/synapse-analytics.html>

QUESTION 96

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 97

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 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
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 99

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 Databricks is an Apache Spark-based analytics platform.	<input type="radio"/>	<input type="radio"/>
Azure Analysis Services is used for transactional workloads.	<input type="radio"/>	<input type="radio"/>
Azure Data Factory orchestrates data integration workflows.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
Azure Databricks is an Apache Spark-based analytics platform.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Analysis Services is used for transactional workloads.	<input type="radio"/>	<input checked="" type="radio"/>
Azure Data Factory orchestrates data integration workflows.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/databricks/scenarios/what-is-azure-databricks>

<https://docs.microsoft.com/en-us/azure/analysis-services/analysis-services-overview>
<https://docs.microsoft.com/en-us/azure/data-factory/introduction>

QUESTION 100

HOTSPOT

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

Hot Area:

Answer Area

Batch workloads ▼

- 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.

Answer Area:

Answer Area

Batch workloads ▼

- 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 101

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	Answer Area
Azure Data Factory	Service Output data to Parquet format
Azure Data Lake Storage	Service Store data that is in Parquet format
Azure SQL Database	Service Persist a tabular representation of data that is stored in Parquet format
Azure Synapse Analytics	

Correct Answer:

Services	Answer Area
	Azure Data Factory Output data to Parquet format
	Azure Data Lake Storage Store data that is in Parquet format
Azure SQL Database	
	Azure Synapse Analytics Persist a tabular representation of data that is stored in Parquet format

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 102

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

Correct Answer:

Services	Answer Area
Azure Blob storage	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
Azure Table storage	

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 103

Which database transaction property ensures that transactional changes to a database are preserved during unexpected operating system restarts?

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

Correct Answer: B

Section:

QUESTION 104

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
Processing salary payments once a month is an example of a batch workload.	<input type="radio"/>	<input type="radio"/>
A wind turbine that sends 50 sensor readings per second is an example of a streaming workload.	<input type="radio"/>	<input type="radio"/>
A home electricity meter that sends readings once a day to an energy provider is an example of a streaming workload.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Statements	Yes	No
Processing salary payments once a month is an example of a batch workload.	<input checked="" type="radio"/>	<input type="radio"/>
A wind turbine that sends 50 sensor readings per second is an example of a streaming workload.	<input checked="" type="radio"/>	<input type="radio"/>
A home electricity meter that sends readings once a day to an energy provider is an example of a streaming workload.	<input type="radio"/>	<input checked="" type="radio"/>

Section:

Explanation:

QUESTION 105

HOTSPOT

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

Hot Area:

Answer Area

In Azure Data Factory, you can use

▼
a control flow
a dataset
a linked service
an integration runtime

 to orchestrate pipeline activities that depend on the output of other pipeline activities.

Answer Area:

Answer Area

In Azure Data Factory, you can use

▼
a control flow
a dataset
a linked service
an integration runtime

 to orchestrate pipeline activities that depend on the output of other pipeline activities.

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/frequently-asked-questions>

QUESTION 106

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 Microsoft Power BI dashboard is associated with a single workspace.	<input type="radio"/>	<input type="radio"/>
A Microsoft Power BI dashboard can only display visualizations from a single dataset.	<input type="radio"/>	<input type="radio"/>
A Microsoft Power BI dashboard can display visualizations from a Microsoft Excel workbook.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
A Microsoft Power BI dashboard is associated with a single workspace.	<input type="radio"/>	<input checked="" type="radio"/>
A Microsoft Power BI dashboard can only display visualizations from a single dataset.	<input type="radio"/>	<input checked="" type="radio"/>
A Microsoft Power BI dashboard can display visualizations from a Microsoft Excel workbook.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/power-bi/connect-data/service-datasets-across-workspaces>

<https://docs.microsoft.com/en-us/power-bi/consumer/end-user-dashboards><https://powerbi.microsoft.com/en-us/excel-and-power-bi/>

QUESTION 107

HOTSPOT

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

Hot Area:

Answer Area

In a data warehousing workload, data

- from a single source is distributed to multiple locations
- from multiple sources is combined in a single location
- is added to a queue for multiple systems to process
- is used to train machine learning models

Answer Area:

Answer Area

In a data warehousing workload, data

- from a single source is distributed to multiple locations
- from multiple sources is combined in a single location
- is added to a queue for multiple systems to process
- is used to train machine learning models

Section:

Explanation:

Note: The data warehouse workload encompasses:

The entire process of loading data into the warehouse

Performing data warehouse analysis and reporting

Managing data in the data warehouse

Exporting data from the data warehouse

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-workload-management>

QUESTION 108

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: <input type="text"/>
The CRM system	Load: <input type="text"/>
The data warehouse	Transform: <input type="text"/>

Correct Answer:

Locations	Answer Area
<input type="text"/>	Extract: <input type="text" value="The CRM system"/>
<input type="text"/>	Load: <input type="text" value="The data warehouse"/>
<input type="text"/>	Transform: <input type="text" value="A standalone data analysis tool"/>

Section:

Explanation:

Box 1: The CRM system

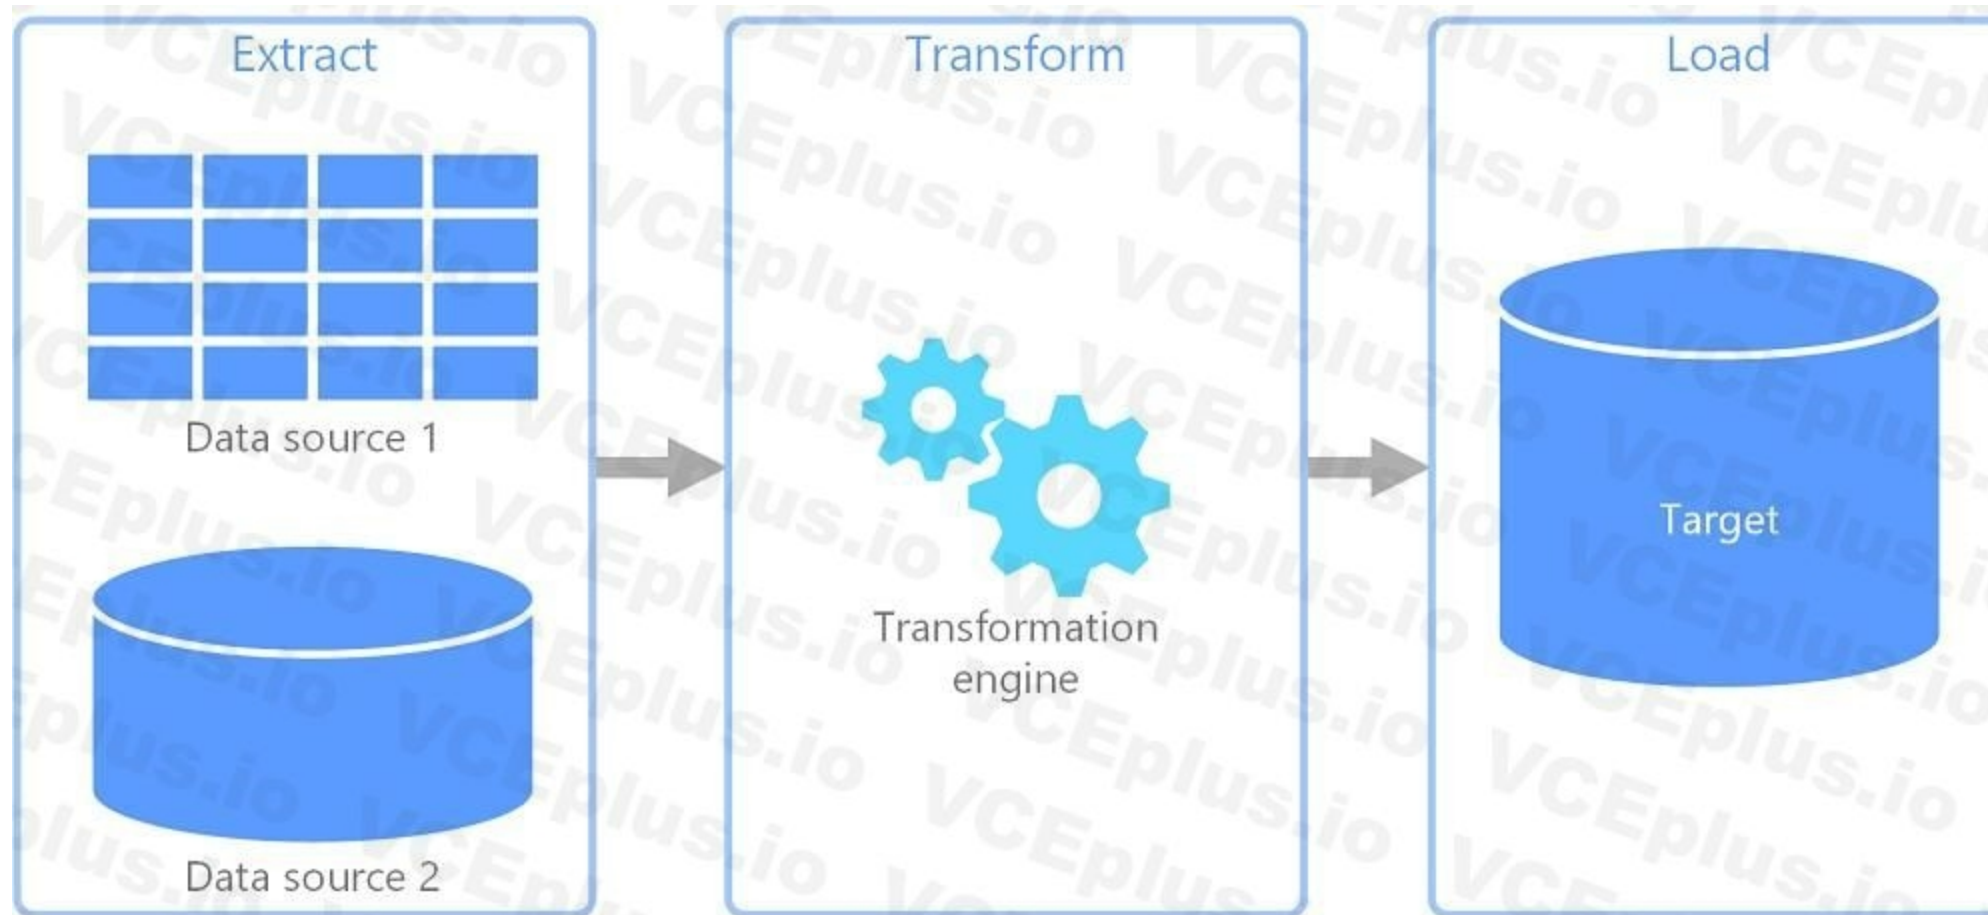
Data is extracted from the CRM system.

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 109

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?
Diagnostic	Provide specific guidance on how to allocate current budget to buy different inventory items.
Descriptive	Detect images of the company's logo included in PDF documents
Predictive	
Prescriptive	

Correct Answer:

Analytics Types	Answer Area
	Why did sales increase last month?
	Provide specific guidance on how to allocate current budget to buy different inventory items.
Descriptive	Detect images of the company's logo included in PDF documents
Predictive	

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 110

DRAG DROP

Match the job roles to the appropriate tasks.

To answer, drag the appropriate job 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 roles	Answer Area
<div>Data analyst</div>	<div></div> Create a pipeline to transfer data from a database to a data warehouse.
<div>Data architect</div>	<div></div> Create a transactional database.
<div>Data engineer</div>	<div></div> Create a chart that shows sales transactions by week.
<div>Database administrator</div>	

Correct Answer:

Job roles	Answer Area
	<div>Data engineer</div> Create a pipeline to transfer data from a database to a data warehouse.
	<div>Data architect</div> Create a transactional database.
<div>Database administrator</div>	<div>Data analyst</div> Create a chart that shows sales transactions by week.

Section:

Explanation:

QUESTION 111

DRAG DROP

Match the database normalization 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:

Terms

Attributes

Entities

Foreign keys

Primary keys

Schemas

Answer Area

Correct Answer:

Terms

Primary keys

Schemas

Answer Area

Section:
Explanation:

www.VCEplus.io