



CompTIA

CertyIQ

Premium exam material

Get certification quickly with the CertyIQ Premium exam material.

Everything you need to prepare, learn & pass your certification exam easily. Lifetime free updates

First attempt guaranteed success.

<https://www.CertyIQ.com>

About CertyIQ

We here at CertyIQ eventually got enough of the industry's greedy exam paid for. Our team of IT professionals comes with years of experience in the IT industry Prior to training CertyIQ we worked in test areas where we observed the horrors of the paywall exam preparation system.

The misuse of the preparation system has left our team disillusioned. And for that reason, we decided it was time to make a difference. We had to make In this way, CertyIQ was created to provide quality materials without stealing from everyday people who are trying to make a living.

Doubt Support

We have developed a very scalable solution using which we are able to solve 400+ doubts every single day with an average rating of 4.8 out of 5.

<https://www.certyiq.com>

[Mail us on - certyiqofficial@gmail.com](mailto:certyiqofficial@gmail.com)



Lifetime Free Updates

We provide lifetime free updates to our customers. To make life easier for our valued customers and fulfill their needs



Free Exam PDF

You are sure to pass the exam completely free of charge



Money Back Guarantee

We Provide 100% money back guarantee to our customer in case of any failure

John

October 19, 2022



Thanks you so much for your help. I scored 972 in my exam today. More than 90% were from your PDFs!

Dana

September 04, 2022



Thanks a lot for this updated AZ-900 Q&A. I just passed my exam and got 974, I followed both of your Az-900 videos and the 6 PDF, the PDFs are very much valid, all answers are correct. Could you please create a similar video/PDF for DP900, your content/PDF's is really awesome. The team did a really good job. Thank You 😊.

Ahamed Shibly

2 months ago



Customer support is really fast and helpful, I just finished my exam and this video along with the 6 PDF helped me pass! Definitely recommend getting the PDFs. Thank you!

October 22, 2022



Passed my exam today with 891 marks. Out of 52 questions, 51 were from certyiq PDFs including Contoso case study. Thank You certyiq team!

Henry Rome

2 months ago



These questions are real and 100 % valid. Thank you so much for your efforts, also your 4 PDFs are awesome, I passed the DP900 exam on 1 Sept. With 968 marks. Thanks a lot, buddy!

Esmaria

2 months ago



Simple easy to understand explanations. To anyone out there wanting to write AZ900, I highly recommend 6 PDF's. Thank you so much, appreciate all your hard work in having such great content. Passed my exam Today - 3 September with 942 score.

Microsoft

(DP-900)

Microsoft Azure Data Fundamentals

Total: **302 Questions**

Link: <https://certyiq.com/papers/microsoft/dp-900>

Question: 1

HOTSPOT -

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

Hot Area:

Answer Area

Descriptive analytics tells you

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

Answer:**Answer Area**

Descriptive analytics tells you

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

Explanation:

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

Reference:

<https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitive-analytics/>**Question: 2**

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements

Yes

No

Normalization involves eliminating relationships between database tables.

☐☐

Normalizing a database reduces data redundancy.

☐☐

Normalization improves data integrity.

☐☐

Answer:

Answer Area

Statements

Yes

No

Normalization involves eliminating relationships between database tables.

☐☒

Normalizing a database reduces data redundancy.

☒☐

Normalization improves data integrity.

☒☐

Explanation:

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

Reference:

<https://docs.microsoft.com/en-us/office/troubleshoot/access/database-normalization-description#:~:text=Normalization%20is%20the%20process%20of,eliminating%20redundancy%20and%20inconsistent%20dependency>

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

Question: 3

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

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

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

Answer:

Answer Area

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

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

Explanation:

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

Reference:

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

Question: 4

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

In batch processing,

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

Answer:

Answer Area

In batch processing,

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

Explanation:

When to use batch processing.

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

Reference:

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

Question: 5

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

Transcribing audio files is an example of

- cognitive
- descriptive
- predictive
- prescriptive

analytics.

Answer:

Answer Area

Transcribing audio files is an example of

- cognitive
- descriptive
- predictive
- prescriptive

analytics.

Explanation:

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

Reference:

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

Question: 6

CertyIQ

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	<input type="text"/>	Why did sales increase last month?
Diagnostic	<input type="text"/>	How do I allocate my budget to buy different inventory items?
Descriptive	<input type="text"/>	Which people are mentioned in a company's business documents?
Predictive		
Prescriptive		

Answer:

Analytics Types

Answer Area

Cognitive	Diagnostic	Why did sales increase last month?
Diagnostic	Predictive	How do I allocate my budget to buy different inventory items?
Descriptive	Descriptive	Which people are mentioned in a company's business documents?
Predictive		
Prescriptive		

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.

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

predictive variables from the past to analyze and identify the likelihood of an unknown future outcome

Box 3: Descriptive -

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



Reference:

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

Question: 7

CertyIQ

HOTSPOT -

You have the following JSON document.


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

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the JSON document.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Customer is [answer choice].

	▼
a nested array	
a nested object	
a root object	

Address is [answer choice].

	▼
a nested array	
a nested object	
a root object	

Social media is [answer choice].

	▼
a nested array	
a nested object	
a root object	

Answer:

Answer Area

Customer is [answer choice].

	▼
a nested array	
a nested object	
a root object	

Address is [answer choice].

	▼
a nested array	
a nested object	
a root object	

Social media is [answer choice].

	▼
a nested array	
a nested object	
a root object	

Explanation:

Reference:

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

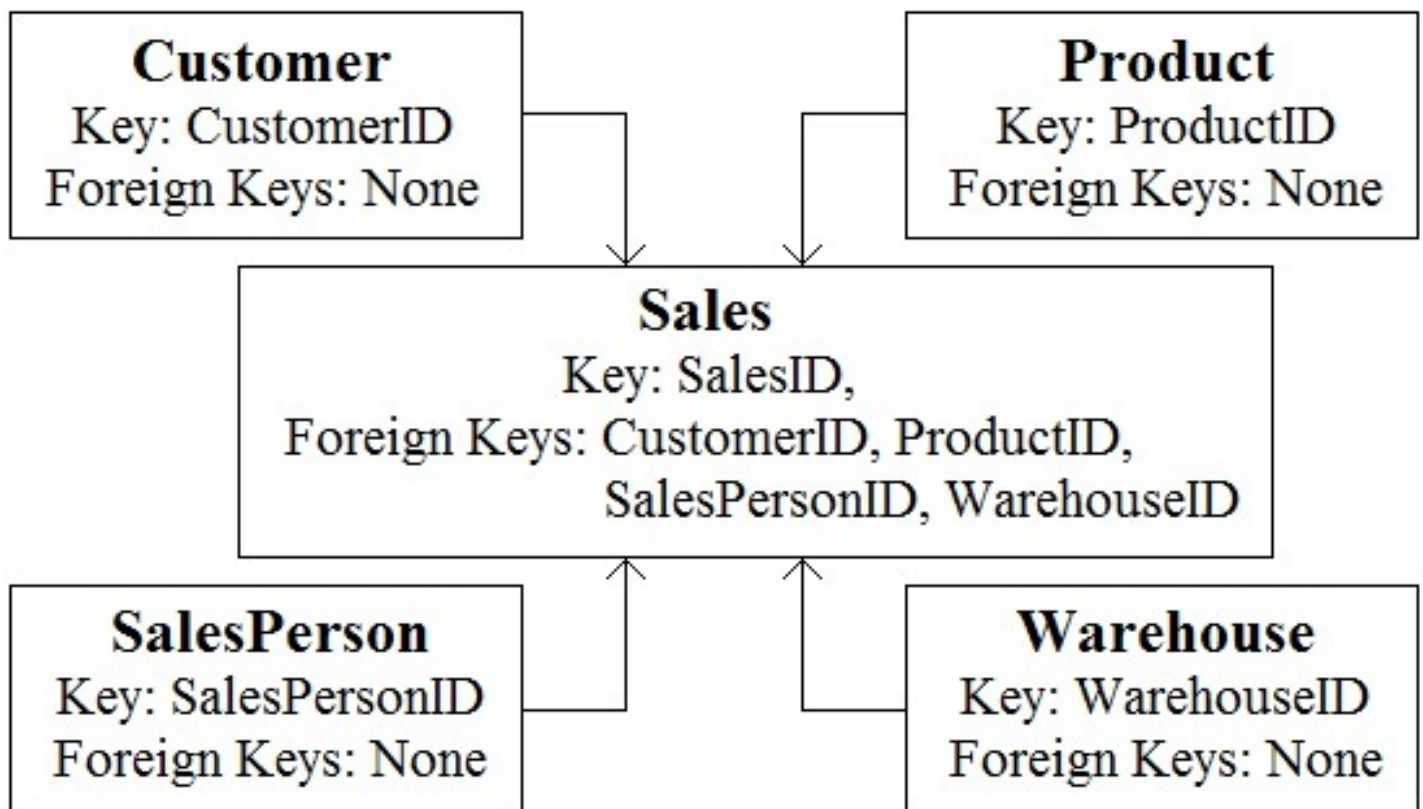
https://www.w3schools.com/js/js_json_objects.asp

Question: 8

CertyIQ

HOTSPOT -

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



Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point

Hot Area:

Answer Area

The data model is a [answer choice].

	▼
transactional model	
star schema	
snowflake schema	

Customer is a [answer choice] table.

	▼
fact	
dimension	
bridge	

Answer:

Answer Area

The data model is a [answer choice].

transactional model
star schema
snowflake schema

Customer is a [answer choice] table.

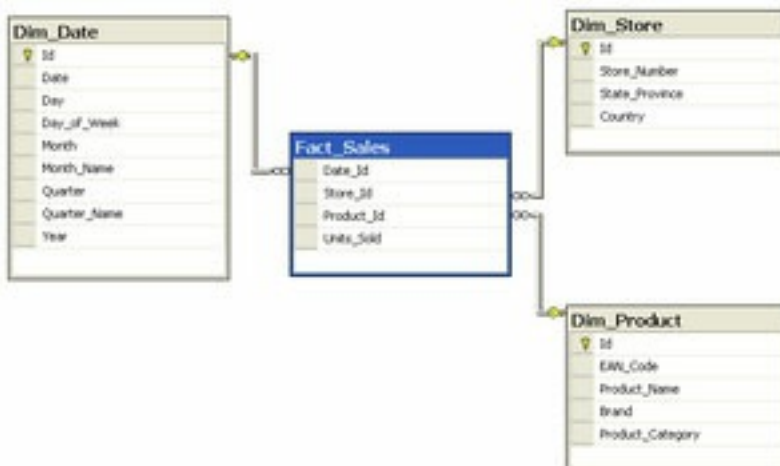
fact
dimension
bridge

Explanation:

Box 1: star schema -

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

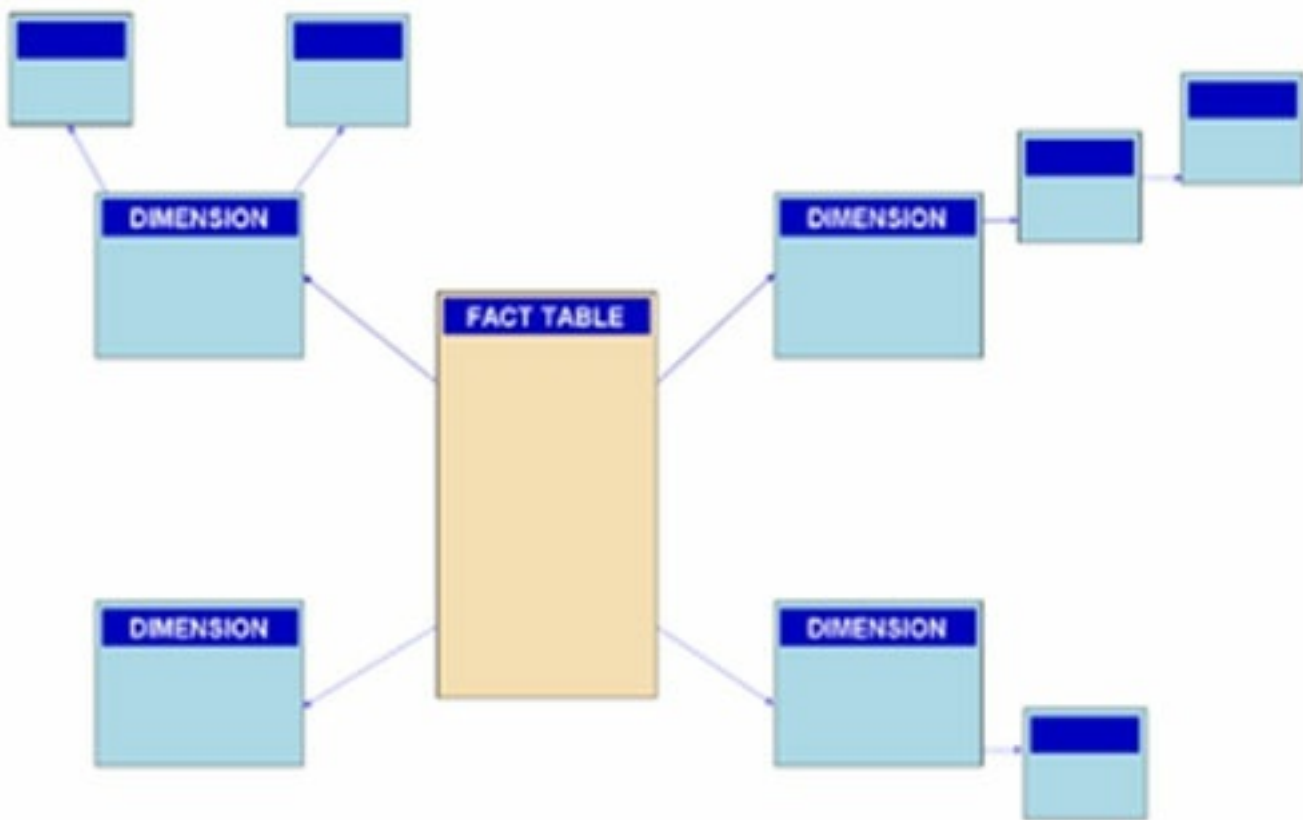
Example:



Incorrect Answers:

The data in the question is not normalized.

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



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

Box 2: dimension -

The star schema consists of one or more fact tables referencing any number of dimension tables.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-overview> https://en.wikipedia.org/wiki/Star_schema https://en.wikipedia.org/wiki/Snowflake_schema <https://azure.microsoft.com/en-us/blog/data-models-within-azure-analysis-services-and-power-bi/>

Question: 9

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

The massively parallel processing (MPP) engine of Azure

Synapse Analytics

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

Answer:

Answer Area

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

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

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/massively-parallel-processing-mpp-architecture>

Question: 10

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

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

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

Answer:

Answer Area

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

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

Explanation:

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

Reference:

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

Question: 11

CertyIQ

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:

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	

Explanation:

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

Note: Relational databases are optimized for writes. They are optimized for consistency and availability.

Advantages of relational databases include simplicity, ease of data retrieval, data integrity, and flexibility.

Incorrect Answers:

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

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

Reference:

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

Question: 12

CertyIQ

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:

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

Explanation:

Batch processing can output data to a file store:

Correct Answer: Yes

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

Batch processing can output data to a relational database:

Correct Answer: Yes

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

Batch processing can output data to a NoSQL database:

Correct Answer: Yes

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

Reference:

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

Question: 13

CertyIQ

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

An in-memory data integration tool

The CRM system

The data warehouse

Answer Area

Extract:

Location

Load:

Location

Transform:

Location

Answer:

Locations

An in-memory data integration tool

The CRM system

The data warehouse

Answer Area

Extract:

The CRM system

Load:

The data warehouse

Transform:

The data warehouse

Explanation:

Extract: The CRM System

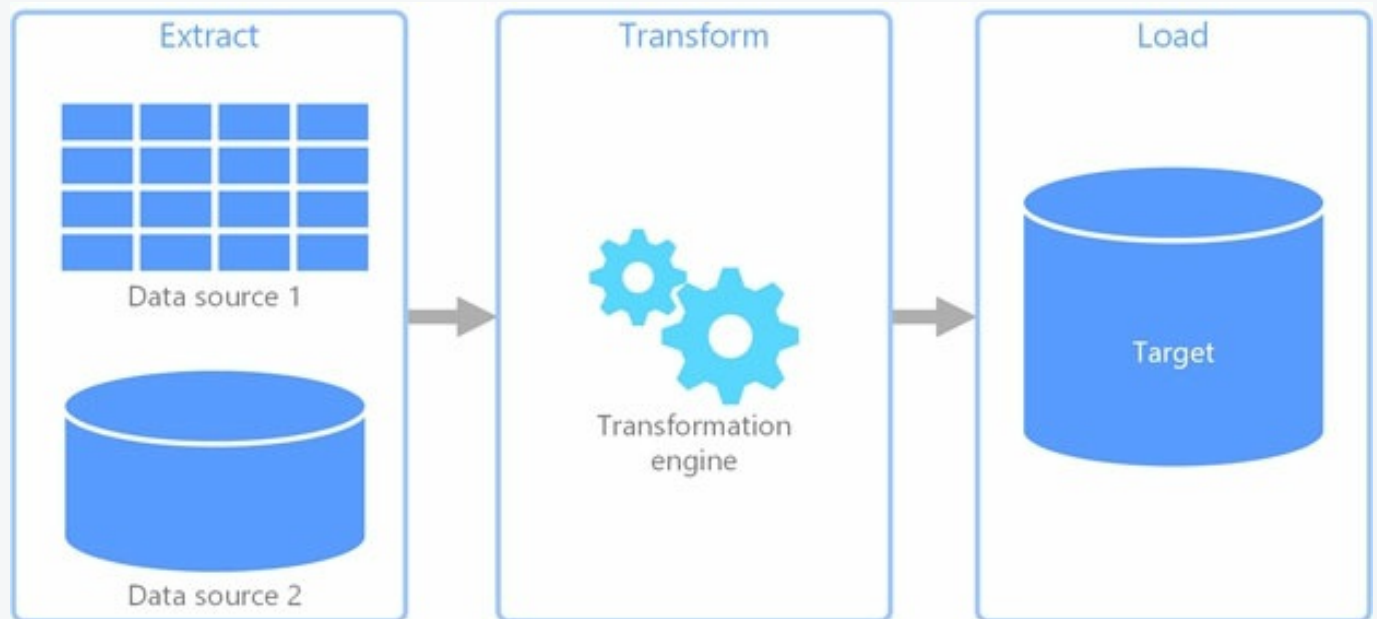
Explanation: In the ELT process, data is extracted directly from the source system, which in this case is the CRM system.

Load: The Data Warehouse

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

Transform: The Data Warehouse

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



Reference:

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

Question: 14

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

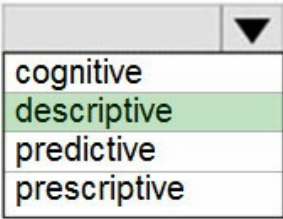
A visualization that shows a university's current student enrollment versus the maximum capacity is an example of

	▼
cognitive	
descriptive	
predictive	
prescriptive	

analytics.

Answer:

Answer Area

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

cognitive
descriptive
predictive
prescriptive

Explanation:

Generally speaking, data analytics comes in four types (Figure 1):

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

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

Predictive, to answer the question: What will happen?

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

Reference:

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

Question: 15

CertyIQ

DRAG DROP -

Match the types of visualizations to the appropriate descriptions.

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

NOTE: Each correct match is worth one point.

Select and Place:

Visualizations

Key influencer

Scatter

Treemap

Answer Area

Visualization

A chart of colored, nested rectangles that displays individual data points represented by the size and color of a relative rectangle.

Visualization

A chart that displays the major contributors of a selected result or value.

Visualization

A chart that shows the relationship between two numerical values.

Answer:

Visualizations

Key influencer

Scatter

Treemap

Answer Area

Treemap

A chart of colored, nested rectangles that displays individual data points represented by the size and color of a relative rectangle.

Key influencer

A chart that displays the major contributors of a selected result or value.

Scatter

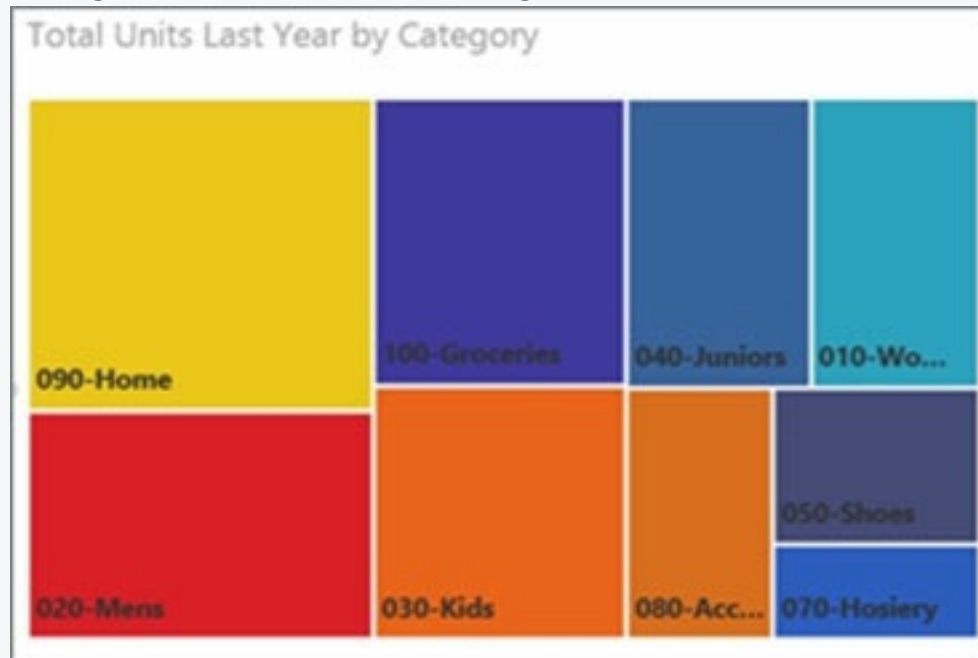
A chart that shows the relationship between two numerical values.

Explanation:

Box 1: Tree map -

Treemaps are charts of colored rectangles, with size representing value. They can be hierarchical, with

rectangles nested within the main rectangles.



Box 2: Key influencer -

A key influencer chart displays the major contributors to a selected result or value.

Box 3: Scatter -

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

Question: 16

CertyIQ

You need to create an Azure Storage account.

Data in the account must replicate outside the Azure region automatically.

Which two types of replication can you use for the storage account? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

Answer: BD

Explanation:

D: Azure Storage offers two options for copying your data to a secondary region:

- ⇒ Geo-redundant storage (GRS)
- ⇒ Geo-zone-redundant storage (GZRS)

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

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy#redundancy-in-a-secondary-region>

Question: 17

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than infrastructure as a service (IaaS) database offerings.	<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:

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"/>

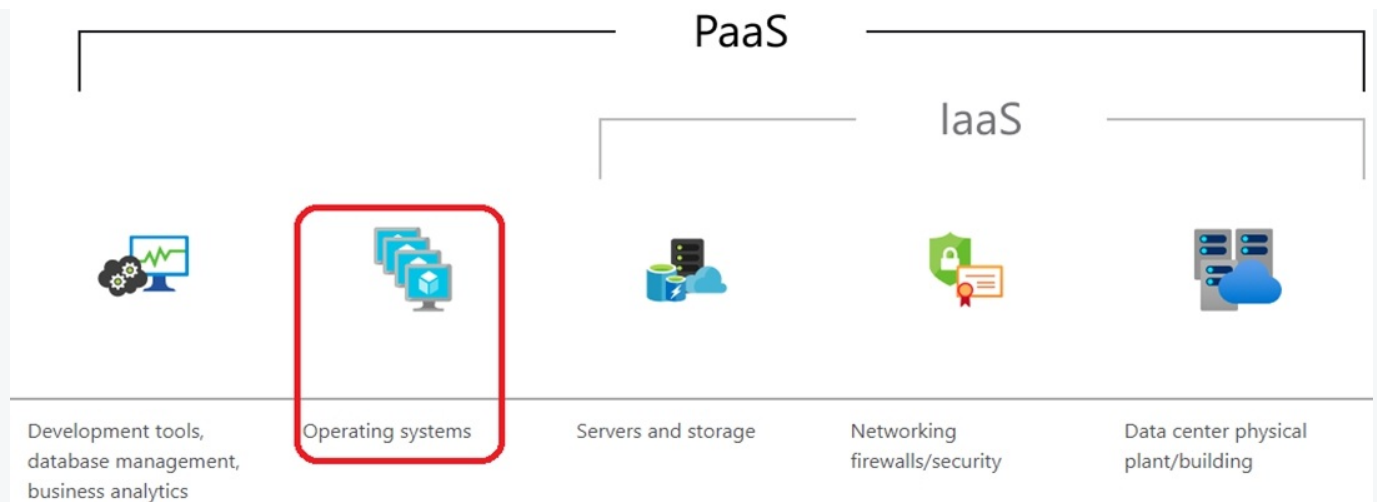
Explanation:

Box 1: Yes -

Like IaaS, PaaS includes infrastructure " servers, storage, and networking " but also middleware, development tools, business intelligence (BI) services, database management systems, and more. PaaS is designed to support the complete web application lifecycle: building, testing, deploying, managing, and updating. PaaS allows you to avoid the expense and complexity of buying and managing software licenses, the underlying application infrastructure and middleware, container orchestrators such as Kubernetes, or the development tools and other resources

Box 2: No -

You manage the applications and services you develop, and the cloud service provider typically manages everything else.



Box 3: No -

There really is no way to pause / stop billing for your Azure SQL Database.

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

Reference:

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

Question: 18

CertyIQ

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

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

Answer: C

Explanation:

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

BULK INSERT -

-
- ⇒ DELETE
- ⇒ INSERT
- ⇒ SELECT
- ⇒ UPDATE
- ⇒ MERGE

Reference:

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

Question: 19

CertyIQ

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

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

Answer: B

Explanation:

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

Reference:

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

Question: 20

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

A key/value data store is optimized for

	▼
enforcing constraints	
simple lookups	
table joins	
transactions	

Answer:

Answer Area

A key/value data store is optimized for

	▼
enforcing constraints	
simple lookups	
table joins	
transactions	

Explanation:

Box 1: simple lookups -

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

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

Reference:

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

Question: 21

CertyIQ

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

Answer:

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

Explanation:

1. Image files = Azure Blob storage

2. Relationship between employees = Azure Cosmos DB Gremlin API

3. Key/value pairs = Azure Table Storage

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

Azure Table Storage is a NoSQL store that efficiently handles structured data in the form of key/value pairs.

Azure Cosmos DB Gremlin API is used for graph-based data, which is ideal for representing relationships, such as the hierarchy or connections between employees.

Reference:

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

Question: 22

DRAG DROP -

Match the Azure Data Lake Storage Gen2 terms to the appropriate levels in the hierarchy.

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

NOTE: Each correct match is worth one point.

Select and Place:

Terms

Azure Storage account

File share

Container

Answer Area

Azure Resource Group

Term

Term

Folders

Files

Answer:

Terms

Azure Storage account

File share

Container

Answer Area

Azure Resource Group

Azure Storage account

Container

Folders

Files

Explanation:

Box 1: Azure Storage account -

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

Box 2: Container -

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

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

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

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

Reference:

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

Question: 23

CertyIQ

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

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

Answer: BD

Explanation:

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

Reference:

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

Question: 24

CertyIQ

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

Answer:

Components

Dataset

Linked service

Mapping data flow

Pipeline

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

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: 25

CertyIQ

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

Batch

Streaming

Answer Area

Workload type

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

Workload type

Thousands of data sets per second for online purchases will be loaded into a data warehouse in real time

Workload type

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

Answer:

Workload Types

Batch

Streaming

Answer Area

Batch

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

Streaming

Thousands of data sets per second for online purchases will be loaded into a data warehouse in real time

Batch

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

Explanation:

Box 1: Batch -

Batch processing refers to the processing of blocks of data that have already been stored over a period of time.

Box 2: Streaming -

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

Box 3: Batch -

Reference:

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

Question: 26

CertyIQ

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:

Answer:

Locations

Answer Area

A standalone data analysis tool

Extract:

The CRM system

The CRM system

Load:

The data warehouse

The data warehouse

Transform:

The data warehouse

Explanation:

Box 1: The CRM system -

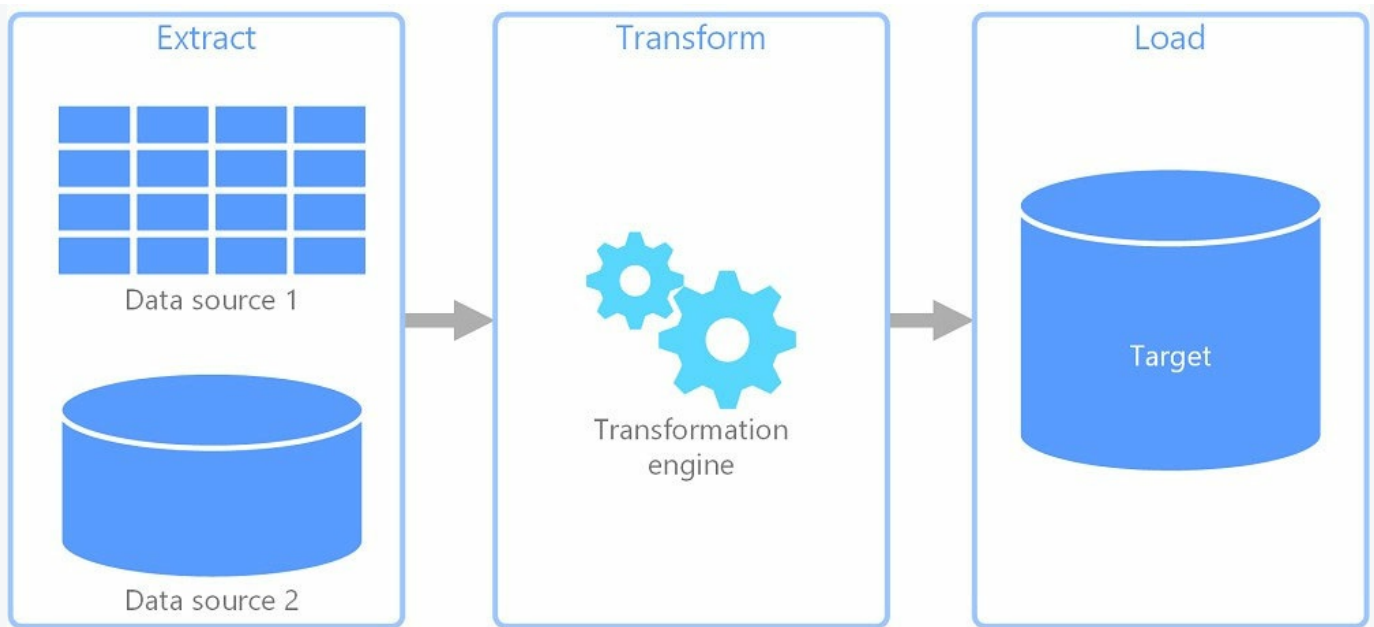
Data is extracted from the CRM system.

Box 2: The data warehouse -

Data is loaded to the data warehouse.

Box 3: The data warehouse -

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




Reference:

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

Question: 27

CertyIQ

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



Which type of analysis is shown in the fourth quarter?

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

Answer: A

Explanation:

Predictive, to answer the question: What will happen?



Descriptive
What's happening?



Diagnostic
Why's happening?



Predictive
What will happen?



Prescriptive
What actions should we take?

Reference:

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

Question: 28

CertyIQ

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

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

Answer: D

Explanation:

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

BULK INSERT -

DELETE -

INSERT -

SELECT -

UPDATE -

MERGE -

There's 3 type of Relational data SQL:

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

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

DCL (Data Control Language) REVOKE & GRANT.

TCL (Transaccional Control Language) COMMIT & ROLLBACK

Reference:

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

Question: 29

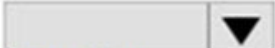
CertyIQ

HOTSPOT -

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

Hot Area:

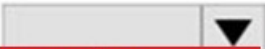
Answer Area

Creating closed caption text for audio files is an example of  analytics.

cognitive
descriptive
predictive
prescriptive

Answer:

Answer Area

Creating closed caption text for audio files is an example of  analytics.

cognitive
descriptive
predictive
prescriptive

Explanation:

Correct answer is Cognitive.

Reference:

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

Question: 30

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

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

▼
cognitive
descriptive
predictive
prescriptive

Answer:

Answer Area

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

▼
cognitive
descriptive
predictive
prescriptive

Explanation:

Generally speaking, data analytics comes in four types:

1. Descriptive, to answer the question: What's happening?
2. Diagnostic, to answer the question: Why's happening?
3. Predictive, to answer the question: What will happen?
4. Prescriptive, to answer the question: What actions should we take?



Descriptive
What's happening?



Diagnostic
Why's happening?



Predictive
What will happen?



Prescriptive
What actions should we take?

Reference:

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

Question: 31

HOTSPOT -

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

Hot Area:

Answer Area

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

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

Answer:**Answer Area**

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

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

Explanation:

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

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>**Question: 32**

You need to create an Azure resource to store data in Azure Table storage.

Which command should you run?

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

Answer: B**Explanation:**

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

Question: 33**CertyIQ**

HOTSPOT -

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

Hot Area:

Answer Area

The Azure Cosmos DB

Core (SQL)
Gremlin
MongoDB
Table

 API enables the use of SELECT statements to

retrieve documents from Azure Cosmos DB.

Answer:**Answer Area**

The Azure Cosmos DB

Core (SQL)
Gremlin
MongoDB
Table

 API enables the use of SELECT statements to

retrieve documents from Azure Cosmos DB.

Explanation:

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

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api>**Question: 34****CertyIQ**

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

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

Answer: B**Explanation:**

"Modify a View" = Alter View

There's 3 type of Relational data SQL:

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

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

DCL (Data Control Language) REVOKE & GRANT.

TCL (Transactional Control Language) COMMIT & ROLLBACK

Question: 35

CertyIQ

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Answer:

Answer Area

Statements	Yes	No
Extract, transform, and load (ETL) can reduce the transfer of sensitive data to destination systems.	<input checked="" type="radio"/>	<input type="radio"/>
Extract, load, and transform (ELT) transforms data by using a compute resource independent of the source system and destination system.	<input type="radio"/>	<input checked="" type="radio"/>
Extract, load, and transform (ELT) minimizes the time it takes to copy large volumes of data to destination systems.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

"Extract, transform, and load (ETL) can reduce the transfer of sensitive data to destination systems." → **Yes**

ETL processes data before loading it into the destination system. This means that sensitive or unnecessary data can be filtered out during the transformation step, reducing exposure and improving security.

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

"Extract, load, and transform (ELT) transforms data by using a compute resource independent of the source system and destination system." → **No**

In ELT, data is first loaded into the destination system (usually a data warehouse like Snowflake, BigQuery, or Redshift), and then transformations occur within the destination system itself.

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

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

"Extract, load, and transform (ELT) minimizes the time it takes to copy large volumes of data to destination systems." → **Yes**

ELT is optimized for handling large datasets because it loads data in bulk first and then processes it within the destination system.

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

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

HOTSPOT -

You plan to deploy a PostgreSQL database to Azure.

Which hosting model corresponds to the available deployment options? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

PostgreSQL on Azure VMs:

▼
Infrastructure as a service (IaaS)
Platform as a service (PaaS)
Software as a service (SaaS)

Azure Database for PostgreSQL:

▼
Infrastructure as a service (IaaS)
Platform as a service (PaaS)
Software as a service (SaaS)

Answer:

Answer Area

PostgreSQL on Azure VMs:

▼
Infrastructure as a service (IaaS)
Platform as a service (PaaS)
Software as a service (SaaS)

Azure Database for PostgreSQL:

▼
Infrastructure as a service (IaaS)
Platform as a service (PaaS)
Software as a service (SaaS)

Explanation:

PostgreSQL on Azure VMs

1. IaaS (Infrastructure as a Service)

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

Azure Database for PostgreSQL

2. PaaS (Platform as a Service)

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

Reference:

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

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

Question: 37

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

Column family databases
Document databases
Graph databases
Key-value stores

natively support the analysis of relationships between entities.

Answer:

Answer Area

Column family databases
Document databases
Graph databases
Key-value stores

natively support the analysis of relationships between entities.

Explanation:

Graph databases:

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

Why the other options are incorrect:

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

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

do not specialize in managing or querying relationships.

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

Important Tip:

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

Question: 38

CertyIQ

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

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

Answer: B

Explanation:

Reference:

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

Question: 39

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

Relational data is stored in

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

Answer:

Answer Area

Relational data is stored in

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

Question: 40

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

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

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

Answer:

Answer Area

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

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

Explanation:

Answer is B, Clustered indexes.

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

Reference:

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

Question: 41

CertyIQ

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

	▼
a stored procedure.	
a table.	
a view.	
an index.	

Answer:

Answer Area

A block of code that runs in a database is called

	▼
a stored procedure.	
a table.	
a view.	
an index.	

Explanation:

Reference:

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

Question: 42

CertyIQ

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		
Predictive	Detect images of the company's logo included in PDF documents	
Prescriptive		

Answer:

Analytics Types

Answer Area

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

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.

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

Box 3: Cognitive -

Reference:

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

Question: 43

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

	▼
A heap	
A stored procedure	
A view	
An index	

presents content defined by a query.

Answer:

Answer Area

	▼
A heap	
A stored procedure	
A view	
An index	

presents content defined by a query.

Question: 44

CertyIQ

Which type of database is Azure Database for PostgreSQL?

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

Answer: A

Explanation:

Correct answer: A, PAAS

Reference:

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

Question: 45

CertyIQ

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

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

Answer: A

Explanation:

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

Reference:

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

Question: 46

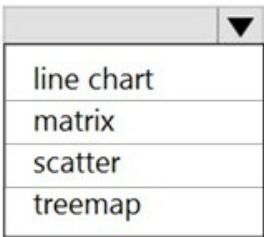
CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

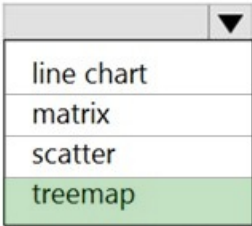
Answer Area

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

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

Answer:

Answer Area

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

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-reporting-services?view=sql-server-ver15>

Question: 47

CertyIQ

What is a characteristic of batch processing?

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

Answer: C

Explanation:

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

Question: 48

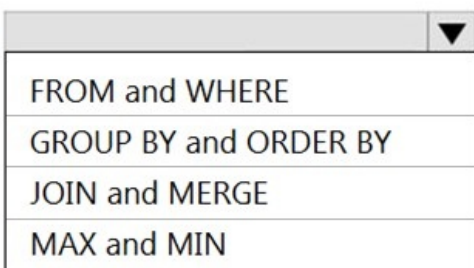
CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

In a SQL statement,  are examples of aggregate functions.

Answer:

Answer Area

In a SQL statement,

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

 are examples of aggregate functions.

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: 49

CertyIQ

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements

Yes

No

A job that processes sales data once daily is an example of a batch workload.

☐☐

A job that calculates a rolling average temperature reading is an example of a streaming workload.

☐☐

A job that calculates average revenue per product for the last month is an example of a streaming workload.

☐☐

Answer:

Answer Area

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

Explanation:

Box 1: Yes -

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

Box 2: Yes -

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

Box 3: No -

Reference:

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

Question: 50

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

In a relational database, each row in a table has

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

Answer:

Answer Area

In a relational database, each row in a table has

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

Explanation:

The same set of columns.

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

Incorrect:

Not: a key value pair.

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

Reference:

<https://www.techopedia.com/definition/4425/database-row>

<https://www.techtarget.com/searchdatamanagement/tip/NoSQL-database-types-explained-Key-value-store>

Question: 51

CertyIQ

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.	<input type="radio"/>	<input type="radio"/>
Analytics systems are more normalized than transactional systems.	<input type="radio"/>	<input type="radio"/>
Normalizing a database results in queries that require more joins.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements	Yes	No
Normalizing a database increases the throughput of writing transactions.	<input type="radio"/>	<input checked="" 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"/>

Explanation:

Box 1: No -

Database normalization is the process of restructuring a relational database in accordance with a series of so-called normal forms in order to reduce data redundancy and improve data integrity.

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

Box 2: No -

Analytics systems are denormalized to increase performance.

Transactional database systems are normalized to increase data consistency.

Box 3: Yes -

Transactional database systems are more normalized and requires more joins.

Reference:

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

Question: 52

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

<div><div></div><div>▼</div></div> <div>Column family databases</div> <div>Document databases</div> <div>Graph databases</div> <div>Key-value stores</div>	natively support the analysis of relationships between entities
--	---

Answer:

Answer Area

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

natively support the analysis of relationships between entities

Explanation:

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

Reference:

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

Question: 53

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A key/value data store is optimized for

	▼
enforcing constraints.	
simple lookups.	
table joins.	

Answer:

Answer Area

A key/value data store is optimized for

	▼
enforcing constraints.	
simple lookups.	
table joins.	

Explanation:

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

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

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

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

Reference:

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

Question: 54

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

In an infrastructure as a service (IaaS) instance of Microsoft SQL Server on Azure, you manage the

	▼
elastic pool	
MySQL server	
PostgreSQL server	
virtual machine	

that hosts SQL Server

Answer:

Answer Area

In an infrastructure as a service (IaaS) instance of Microsoft SQL Server on Azure, you manage the

	▼
elastic pool	
MySQL server	
PostgreSQL server	
virtual machine	

that hosts SQL Server

Explanation:

SQL Server on Azure VM -

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

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

Reference:

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

Question: 55

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

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

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

Answer:

Answer Area

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

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

Explanation:

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

Reference:

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

Question: 56

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A JSON document is an example of

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

Answer:

Answer Area

A JSON document is an example of

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

Explanation:

Semi-structured data -

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

Reference:

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

Question: 57

CertyIQ

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

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

Answer: A

Explanation:

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

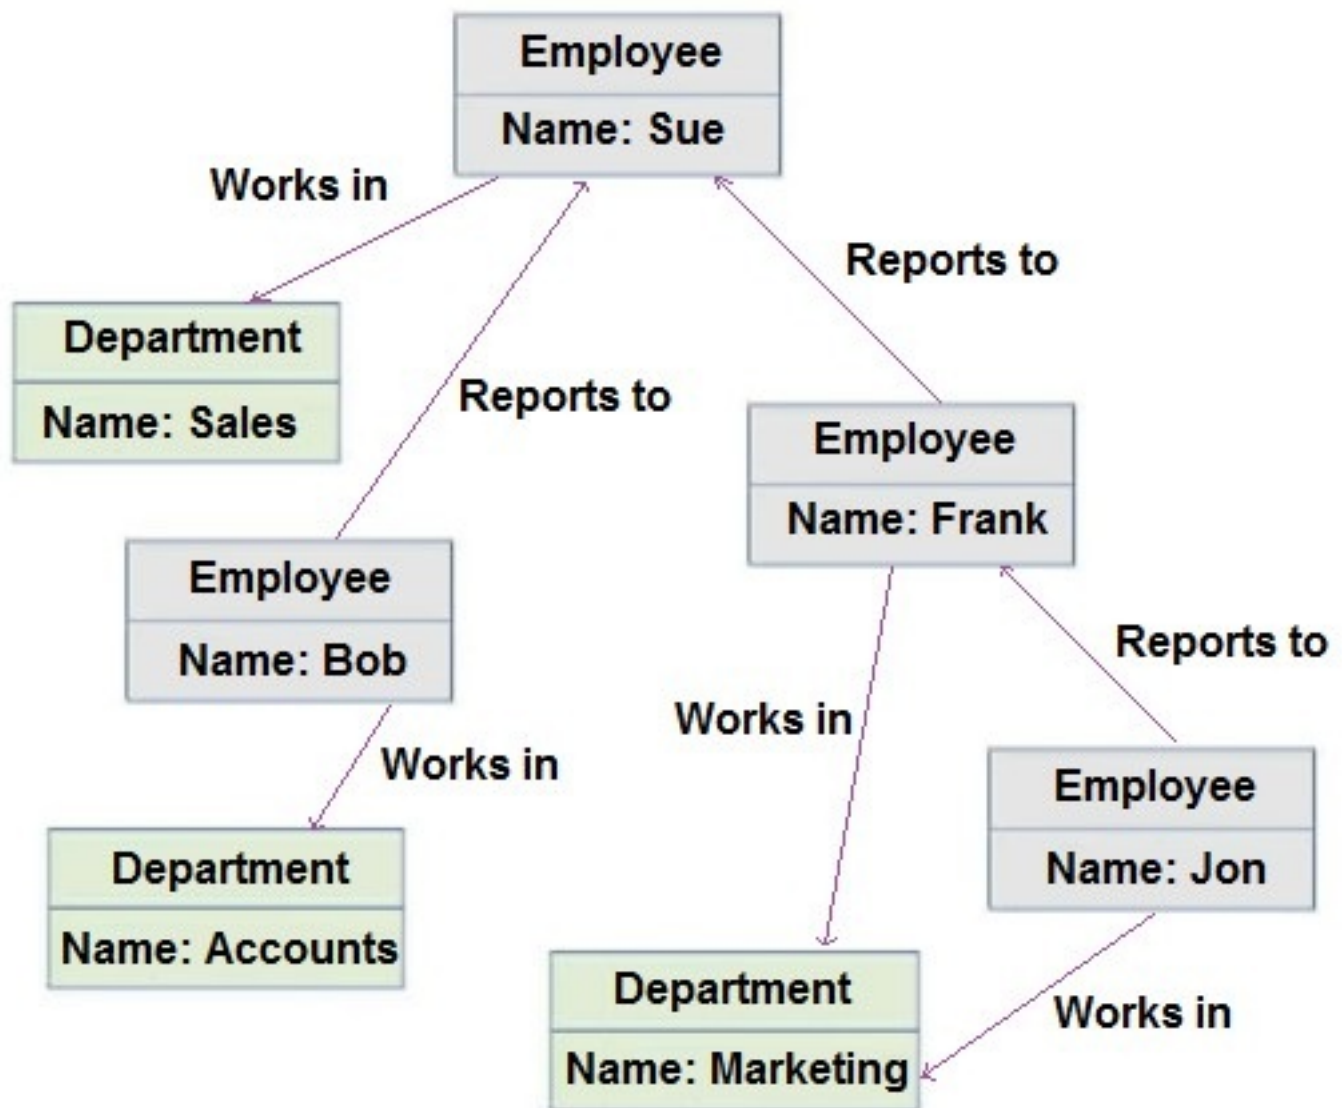
Reference:

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

Question: 58

CertyIQ

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



Which type of data store is this?

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

Answer: C

Explanation:

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

Reference:

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

Question: 59

CertyIQ

Which type of database contains nodes and edges?

- A. graph
- B. key/value

- C. columnar
- D. time series

Answer: A

Explanation:

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

Reference:

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

Question: 60

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

	▼
A data analyst	
A data engineer	
A data scientist	

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

Answer:

Answer Area

	▼
A data analyst	
A data engineer	
A data scientist	

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

Explanation:

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

Incorrect:

- * A data engineer requires an intermediate level understanding of programming to build thorough algorithms along with a mastery of statistics and math.
- * A data scientist needs to be a master of both worlds. Data, stats, and math along with in-depth programming knowledge for Machine Learning and Deep Learning.

Reference:

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

Question: 61

HOTSPOT -

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

Hot Area:

Answer AreaRelational data uses

collections

rows

keys

partitions

Answer:**Answer Area**Relational data uses

collections

rows

keys

partitions

Explanation:

Reference:

<https://teachcomputerscience.com/relational-databases/>**Question: 62**

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

ProductName	Quantity
Product1	100
Product2	129
Product3	176

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

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

Answer: C**Explanation:**

ALTER - add/delete/modify column

UPDATE - update existing row

INSERT - Insert new row

Reference:

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

Question: 63

CertyIQ

Your company needs to implement a relational database in Azure. The solution must minimize ongoing maintenance.

Which Azure service should you use?

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

Answer: B

Explanation:

Reference:

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

Question: 64

CertyIQ

You are writing a set of SQL queries that administrators will use to troubleshoot an Azure SQL database. You need to embed documents and query results into a SQL notebook.

What should you use?

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

Answer: B

Explanation:

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

Reference:

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

Question: 65

CertyIQ

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:

Terms

Index
View
Table

Answer Area

	A database object that holds data
	A database object whose content is defined by a query
	A database object that helps improve the speed of data retrieval

Answer:

Terms

Index
View
Table

Answer Area

Table	A database object that holds data
View	A database object whose content is defined by a query
Index	A database object that helps improve the speed of data retrieval

Explanation:

Reference:

[https://en.wikipedia.org/wiki/Table_\(database\)](https://en.wikipedia.org/wiki/Table_(database))

[https://en.wikipedia.org/wiki/View_\(SQL\)](https://en.wikipedia.org/wiki/View_(SQL))

https://en.wikipedia.org/wiki/Database_index

Question: 66

CertyIQ

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

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

Answer: D

Explanation:

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

Online Transaction Processing (OLTP)

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

core-data-concepts

Reference:

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

Question: 67

CertyIQ

When can you use an Azure Resource Manager template?

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

Answer: A

Explanation:

You can automate deployments and use the practice of infrastructure as code. In code, you define the infrastructure that needs to be deployed

To implement infrastructure as code for your Azure solutions, use Azure Resource Manager templates (ARM templates). The template is a JavaScript Object

Notation (JSON) file that defines the infrastructure and configuration for your project. The template uses declarative syntax, which lets you state what you intend to deploy without having to write the sequence of programming commands to create it. In the template, you specify the resources to deploy and the properties for those resources.

Reference:

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

Question: 68

CertyIQ

You have an Azure SQL database that you access directly from the Internet.

You recently changed the public IP address of your computer.

After changing the IP address, you can no longer access the database. You can connect to other resources in Azure.

What is a possible cause of the issue?

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

Answer: D

Explanation:

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

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

Reference:

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

Question: 69

CertyIQ

DRAG DROP -

Match the tools to the appropriate descriptions.

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

Select and Place:

Tools	Answer Area
Azure Data Studio	Tool A graphical tool for managing SQL Server or Azure SQL databases that supports access, configuration, management, and administration tasks.
Microsoft SQL Server Data Tools (SSDT)	Tool A lightweight editor that can run on-demand SQL queries and view and save results as text, JSON, or Microsoft Excel files.
Microsoft SQL Server Management Studio (SSMS)	Tool A development tool for building Azure SQL databases, Microsoft SQL Server relational databases, SQL Server Analysis Services (SSAS) data models, SQL Server Integration Services (SSIS) packages, and SQL Server Reporting Services (SSRS) reports.

Answer:

Tools	Answer Area
	Microsoft SQL Server Management Studio (SSMS) A graphical tool for managing SQL Server or Azure SQL databases that supports access, configuration, management, and administration tasks.
	Azure Data Studio A lightweight editor that can run on-demand SQL queries and view and save results as text, JSON, or Microsoft Excel files.
	Microsoft SQL Server Data Tools (SSDT) A development tool for building Azure SQL databases, Microsoft SQL Server relational databases, SQL Server Analysis Services (SSAS) data models, SQL Server Integration Services (SSIS) packages, and SQL Server Reporting Services (SSRS) reports.

Explanation:

Box 1: Microsoft SQL Server Management Studio (SSMS)

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

Box 2: Azure Data Studio -

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

Box 3: Microsoft SQL Server Data Tools (SSDT)

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

Reference:

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

Question: 70

CertyIQ

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
Relational database tables contain columns and rows	<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:

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"/>

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>

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

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

Answer: A

Explanation:

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

Incorrect Answers:

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

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

Reference:

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

Question: 72

CertyIQ

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
Azure SQL Database includes a managed backup service.	<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 Defender.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements	Yes	No
Azure SQL Database includes a 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 Defender.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Box 1: Yes -

Box 2: Yes -

Box 3: Yes -

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

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

Reference:

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

Question: 73

CertyIQ

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:

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"/>

Explanation:

A) **Yes** from <https://docs.microsoft.com/en-us/sql/azure-data-studio/what-is?view=sql-server-ver15>

Use Azure Data Studio if you: Are connecting to a SQL Server 2019 big data cluster

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

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

Reference:

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

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-query-ssms> <https://docs.microsoft.com/en-us/azure/mariadb/connect-workbench>

Question: 74

CertyIQ

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:

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"/>

Explanation:

Reference:

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

Question: 75

CertyIQ

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:

Answer Area

dbo.Products:

	▼
A column	
A database	
A table	
An index	

ProductName:

	▼
A column	
A database	
A table	
An index	

Question: 76

CertyIQ

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.

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:

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"/>

Explanation:

N/N/Y

Can I bring my on-premises or IaaS SQL Server license to Hyperscale? Yes, Azure Hybrid Benefit ...
<https://docs.microsoft.com/en-us/azure/azure-sql/database/service-tier-hyperscale-frequently-asked-questions-faq>

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

Azure Hybrid Benefit allows you to use SQL Server licenses with Software Assurance or qualifying subscription licenses to pay a reduced base rate* for these products and services for SQL Server on Azure:

vCPU-based service tiers of Azure SQL Database (excluding serverless).

Azure SQL Managed Instance.

SQL Server in Azure Virtual Machines.

SQL Server Integration Services.

Reference:

<https://azure.microsoft.com/en-gb/blog/hot-patching-sql-server-engine-in-azure-sql-database/>

<https://azure.microsoft.com/en-us/services/sql-database/#product-overview>

Question: 77

CertyIQ

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

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

Answer: D

Explanation:

Data Definition Language (DDL) statements defines data structures. Use these statements to create, alter, or drop data structures in a database. These statements include:

- ⇒ ALTER
- ⇒ Collations
- ⇒ CREATE
- ⇒ DROP
- ⇒ DISABLE TRIGGER
- ⇒ ENABLE TRIGGER
- ⇒ RENAME
- ⇒ UPDATE STATISTICS
- ⇒ TRUNCATE TABLE
- ⇒ INSERT

Reference:

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

Question: 78

CertyIQ

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
Azure Data Studio can be used to query an Azure SQL database from a device that runs macOS.	<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:

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"/>

Explanation:

Box 1: Yes -

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

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

Box 2: No -

SQL Server Management Studio is for configuring, managing, and administering all components within Microsoft SQL Server, not to create SQL notebooks.

Instead use Azure Data Studio to create SQL notebook.

Box 3: Yes -

You can use the Azure Data Studio to restore databases.

Reference:

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

Question: 79

CertyIQ

You are deploying a software as a service (SaaS) application that requires a relational database for Online Transaction Processing (OLTP).

Which Azure service should you use to support the application?

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

Answer: C

Explanation:

Azure SQL Database is relational database and a managed service.

Incorrect Answers:

A, B: Cosmos DB, HDInsight are non-relational databases.

D: Azure Synapse Analytics is for data warehousing, not for Online Transaction Processing

Reference:

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

Question: 80

CertyIQ

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

NOTE: Each correct selection is worth one point.

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

Answer: AD

Explanation:

A: Azure SQL Database is a fully managed platform as a service (PaaS) database engine that handles most of the database management functions such as upgrading, patching, backups, and monitoring without user involvement.

D: SQL Database delivers predictable performance with multiple resource types, service tiers, and compute sizes. It provides dynamic scalability with no downtime, built-in intelligent optimization, global scalability and availability, and advanced security options. These capabilities allow you to focus on rapid app development and accelerating your time-to-market, rather than on managing virtual machines and infrastructure.

Reference:

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

Question: 81

CertyIQ

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
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 installation of the database engine.	<input type="radio"/>	<input type="radio"/>

Answer:

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 installation of the database engine.	<input type="radio"/>	<input checked="" type="radio"/>

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: 82

CertyIQ

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

GROUP BY

ORDER BY

SELECT

Answer Area

Value

SalesDate, AVG(SalesAmount)

FROM Sales

Value

SalesDate

ORDER BY SalesDate

Answer:

Values

CREATE

GROUP BY

ORDER BY

SELECT

Answer Area

SELECT

SalesDate, AVG(SalesAmount)

FROM Sales

GROUP BY

SalesDate

ORDER BY SalesDate

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: 83

CertyIQ

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

Answer: B

Explanation:

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: 84

CertyIQ

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

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

Answer: D

Explanation:

Data Definition Language (DDL) statements defines data structures. Use these statements to create, alter, or drop data structures in a database. These statements include:

- ALTER
- Collations
- CREATE
- DROP
- DISABLE TRIGGER

- ENABLE TRIGGER
- RENAME
- UPDATE STATISTICS
- TRUNCATE TABLE
- UPDATE

Reference:

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

Question: 85

CertyIQ

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

Answer: C

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: 86

CertyIQ

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:

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. |

Explanation:

Transparent data encryption (TDE) helps protect Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics against the threat of malicious offline activity by encrypting data at rest.

Reference:

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

Question: 87

CertyIQ

You need to ensure that users use multi-factor authentication (MFA) when connecting to an Azure SQL database. Which type of authentication should you use?

- A. service principal authentication
- B. Azure Active Directory (Azure AD) authentication
- C. SQL authentication
- D. certificate authentication

Answer: B

Explanation:

Other authentication types, such as service principal authentication, SQL authentication and certificate authentication, don't allow the use of MFA.

Reference:

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

Question: 88

CertyIQ

What is a benefit of hosting a database on Azure SQL managed instance as compared to an Azure SQL database?

- A. built-in high availability
- B. native support for cross-database queries and transactions
- C. system-initiated automatic backups
- D. support for encryption at rest

Answer: B

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/features-comparison>

Question: 89

CertyIQ

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:

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.

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: 90

CertyIQ

You need to design and model a database by using a graphical tool that supports project-oriented offline database development.

What should you use?

- A. Microsoft SQL Server Data Tools (SSDT)
- B. Microsoft SQL Server Management Studio (SSMS)
- C. Azure Databricks
- D. Azure Data Studio

Answer: A

Explanation:

SQL Server Data Tools (SSDT) is a modern development tool for building SQL Server relational databases, databases in Azure SQL, Analysis Services (AS) data models, Integration Services (IS) packages, and

Reporting Services (RS) reports. With SSDT, you can design and deploy any SQL Server content type with the same ease as you would develop an application in Visual Studio.

Reference:

<https://docs.microsoft.com/en-us/sql/ssdt/project-oriented-offline-database-development?view=sql-server-ver15>

Question: 91

CertyIQ

DRAG DROP -

Match the security components to the appropriate scenarios.

To answer, drag the appropriate component from the column on the left to its scenario 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

Authentication

Firewall

Encryption

Answer Area

Prevent access to an Azure SQL database from another network.

Support Azure Active Directory (Azure AD) sign-ins to an Azure SQL database.

Ensure that sensitive data never appears as plain text in an Azure SQL database.

Answer:

Components

Authentication

Firewall

Encryption

Answer Area

Firewall

Prevent access to an Azure SQL database from another network.

Authentication

Support Azure Active Directory (Azure AD) sign-ins to an Azure SQL database.

Encryption

Ensure that sensitive data never appears as plain text in an Azure SQL database.

Explanation:

Firewall (Example someone from outside cannot log in to your company domain because of different IP address)

Authentication is always used to check who is authorized to log in

Encryption

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-a-windows-firewall-for-database-engine-access?view=sql-server-ver15> <https://docs.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-overview> <https://docs.microsoft.com/en-us/azure/azure-sql/database/always-encrypted-certificate-store-configure>

Question: 92

CertyIQ

You have a transactional application that stores data in an Azure SQL managed instance. When should you implement a read-only database replica?

A. You need to generate reports without affecting the transactional workload.

- B. You need to audit the transactional application.
- C. You need to implement high availability in the event of a regional outage.
- D. You need to improve the recovery point objective (RPO).

Answer: A

Explanation:

Use read-only replicas to offload read-only query workloads.

You should implement a read-only database replica when you need to generate reports without affecting the transactional workload.

A read-only replica is a copy of the database that can be used for reporting or other read-only workloads without affecting the performance of the transactional application. This allows you to offload reporting and other read-only workloads to the replica, which can improve the performance of the transactional application.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/read-scale-out>

Question: 93

CertyIQ

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:

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.	

Question: 94

CertyIQ

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.

	▼
collections	
columns	
keys	
partitions	

Answer:

Answer Area

Relational data uses

	▼
collections	
columns	
keys	
partitions	

 to enforce relationships between different tables.

	▼
collections	
columns	
keys	
partitions	

Explanation:

Again RDBMS concepts primary key, unique key & foreign key concepts

Reference:

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

Question: 95

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

	▼
A heap	
A stored procedure	
A view	
An index	

 is a virtual table that contains content defined by a query.

	▼
A heap	
A stored procedure	
A view	
An index	

Answer:

Answer Area

	▼
A heap	
A stored procedure	
A view	
An index	

is a virtual table that contains content defined by a query.

Explanation:

heap: tables: without a clustered index.

Stored procedure: a sub-routine available to access application in RDBMS

A view: is a table that is being created based on a user query

An index: a data structure that improves the speed of data retrieval operations

The answer is obvious, A View

Reference:

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

Question: 96

CertyIQ

You need to query a table named Products in an Azure SQL database.

Which three requirements must be met to query the table from the internet? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. You must be assigned the Reader role for the resource group that contains the database.
- B. You must have SELECT access to the Products table.
- C. You must have a user in the database.
- D. You must be assigned the Contributor role for the resource group that contains the database.
- E. Your IP address must be allowed to connect to the database.

Answer: BCE

Explanation:

Incorrect Answers:

A, D: Resource group permissions is not required to query an Azure SQL database table.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/getting-started-with-database-engine-permissions?view=sql-server-ver15>

Question: 97

CertyIQ

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

Answer: D

Explanation:

Reference:

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

Question: 98

CertyIQ

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

Answer: D

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: 99

CertyIQ

DRAG DROP -

Match the Azure SQL services to the appropriate use cases.

To answer, drag the appropriate service from the column on the left to its use case 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:

Azure SQL services Answer Area

Azure SQL Database elastic pool		A serverless configuration
Azure SQL Database single database		A database that has the highest compatibility with on-premises Microsoft SQL Server
Azure SQL Managed instance		A low-cost, low-maintenance database

Answer:

Azure SQL services Answer Area

Azure SQL Database elastic pool	Azure SQL Database single database	A serverless configuration
Azure SQL Database single database	Azure SQL Managed instance	A database that has the highest compatibility with on-premises Microsoft SQL Server
Azure SQL Managed instance	Azure SQL Database elastic pool	A low-cost, low-maintenance database

Explanation:

Box 1: Azure SQL Database single database

Serverless is a compute tier for single databases in Azure SQL Database that automatically scales compute based on workload demand and bills for the amount of compute used per second. The serverless compute tier also automatically pauses databases during inactive periods when only storage is billed and automatically resumes databases when activity returns.

Scenarios well suited for serverless compute

Single databases with intermittent, unpredictable usage patterns interspersed with periods of inactivity, and lower average compute utilization over time.

Single databases in the provisioned compute tier that are frequently rescaled and customers who prefer to delegate compute rescaling to the service.

New single databases without usage history where compute sizing is difficult or not possible to estimate prior to deployment in SQL Database.

Box 2: Azure SQL Managed Instance

Azure SQL Managed Instance is the intelligent, scalable cloud database service that combines the broadest SQL Server database engine compatibility with all the benefits of a fully managed and evergreen platform as a service.

Box 3: Azure SQL Database elastic pool

Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price. Elastic pools in SQL Database enable software as a service (SaaS) developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview?view=azuresql> <https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview?view=azuresql> <https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/sql-managed-instance-paas-overview>

Question: 100

CertyIQ

Which T-SQL statement should be used to instruct a database management system to use an index instead of performing a full table scan?

- A. SELECT
- B. WHERE
- C. JOIN

Answer: B

Explanation:

B. WHERE

In T-SQL, the WHERE clause is used to filter records. When an index is available on the columns specified in the WHERE clause, the database management system can use the index to quickly locate the rows that match the condition, avoiding a full table scan.

A WHERE clause is the primary way to instruct the database to use an index because it filters rows, allowing the database engine to leverage indexed columns to quickly locate the required data.

Question: 101

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

	▼
Azure SQL Database	
Azure Synapse Analytics	
Azure Database for MariaDB	
SQL Server on Azure virtual machines	

is a platform as a service (PaaS) relational data service that uses massively parallel processing (MPP).

Answer:

Answer Area

	▼
Azure SQL Database	
Azure Synapse Analytics	
Azure Database for MariaDB	
SQL Server on Azure virtual machines	

is a platform as a service (PaaS) relational data service that uses massively parallel processing (MPP).

Explanation:

Azure Synapse Analytics is an platform as a service (PAAS) that combines data integration, warehousing, and analytics into one solution.

Azure Synapse Analytics offers cloud-based, relational data warehousing services, massively parallel processing (MPP) scale-out technology, and enough computational power to efficiently manage petabytes and petabytes of data.

Incorrect:

* Azure Database for MariaDB is a relational database service based on the open-source MariaDB Server engine. It's a fully managed database as a service offering that can handle mission-critical workloads with predictable performance and dynamic scalability.

* Azure SQL Database is a fully managed platform as a service (PaaS) database engine that handles most of the database management functions such as upgrading, patching, backups, and monitoring without user involvement. Azure SQL Database is always running on the latest stable version of the SQL Server database engine and patched OS with 99.99% availability. PaaS capabilities built into Azure SQL Database enable you to focus on the domain-specific database administration and optimization activities that are critical for your business.

Reference:

<https://cswsolutions.com/blog/posts/2021/august/what-is-azure-synapse-analytics/> <https://www.integrate.io/blog/what-is-azure-synapse-analytics/>

Question: 102**CertyIQ**

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

	▼
A view	
A clustered index	
A stored procedure	
A nonclustered index	

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

Answer:

Answer Area

	▼
A view	
A clustered index	
A stored procedure	
A nonclustered index	

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

Explanation:

A clustered index defines the order in which data is physically stored in a table. Table data can be sorted in only one way, therefore, there can be only one clustered index per table. In SQL Server, the primary key constraint automatically creates a clustered index on that particular column.

Reference:

<https://www.sqlshack.com/what-is-the-difference-between-clustered-and-non-clustered-indexes-in-sql-server>

Question: 103

CertyIQ

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:

Answer Area

dbo.Products:

	▼
A column	
A database	
A table	
An index	

ProductName:

	▼
A column	
A database	
A table	
An index	

Question: 104

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

In a relational database, the

	▼
foreign key	
primary key	
WHERE clause	
nonclustered index	

uniquely identifies each row in a table

Answer:

Answer Area

In a relational database, the

	▼
foreign key	
primary key	
WHERE clause	
nonclustered index	

uniquely identifies each row in a table

Explanation:

A primary key is a special relational database table column (or combination of columns) designated to uniquely identify each table record (row).

Reference:

<https://www.techopedia.com/definition/5547/primary-key>

Question: 105

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

	▼
A heap	
A stored procedure	
A view	
An index	

presents content defined by a query.

Answer:

Answer Area

	▼
A heap	
A stored procedure	
A view	
An index	

presents content defined by a query.

Explanation:

A database view is a subset of a database and is based on a query that runs on one or more database tables.

Reference:

https://www.ibm.com/docs/SSWT9A_7.6.0/com.ibm.mbs.doc/configur/c_views.html

Question: 106

CertyIQ

Which Azure service provides the highest compatibility for databases migrated from Microsoft SQL Server 2019 Enterprise edition?

- A. Azure SQL Database
- B. Azure Database for MySQL
- C. Azure SQL Managed Instance
- D. an Azure SQL Database elastic pool

Answer: C

Explanation:

SQL Managed Instance has near 100% compatibility with the latest SQL Server (Enterprise Edition) database engine, providing a native virtual network (VNet) implementation that addresses common security concerns, and a business model favorable for existing SQL Server customers.

Note: Azure SQL Managed Instance is the intelligent, scalable cloud database service that combines the broadest SQL Server database engine compatibility with all the benefits of a fully managed and evergreen platform as a service.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/sql-managed-instance-paas-overview?view=azuresql>

Question: 107

CertyIQ

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

Answer: D

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: 108

CertyIQ

DRAG DROP -

Match the types of data stores to the appropriate scenarios.

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

NOTE: Each correct match is worth one point.

Select and Place:

Data Store Types

Answer Area

Graph
Object
Key/value

<input type="text"/>	A list of application user names and, for each, the default language
<input type="text"/>	A collection of medical images and their associated metadata.
<input type="text"/>	Employee data that shows the relationships between employees

Answer:

Data Store Types

Graph

Object

Key/value

Answer Area

Key/value

Object

Graph

A list of application user names and, for each, the default language

A collection of medical images and their associated metadata.

Employee data that shows the relationships between employees

Explanation:

Box 1: Key/value -

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

Box 2: Object -

Object storage is optimized for storing and retrieving large binary objects (images, files, video and audio streams, large application data objects and documents, virtual machine disk images).

Box 3: Graph -

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

Reference:

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

Question: 109

CertyIQ

You have an Azure Cosmos DB account that uses the Core (SQL) API.

Which two settings can you configure at the container level? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. the throughput
- B. the read region
- C. the partition key
- D. the API

Answer: AC

Explanation:

Reference:

<https://www.sqlshack.com/start-your-journey-with-azure-cosmos-db/>

Question: 110

CertyIQ

Your company is designing a data store that will contain student data. The data has the following format.

StudentNumber	StudentInformation
7634634	First name: Ben Last: Smith Preferred Name: Benjamin
7634634	First Name: Dominik Last Name: Paiha Email Address: dpaiha@contoso.com MCP ID: 931817
7634636	First Name: Reshma Last Name: Patel Phone number: 514-555-1101
7634637	First Name: Yun-Feng Last Name: Peng

Which type of data store should you use?

- A. graph
- B. key/value
- C. object
- D. columnar

Answer: D

Explanation:

Correct answer is D: columnar.

Question: 111

CertyIQ

Which storage solution supports role-based access control (RBAC) at the file and folder level?

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

Answer: B

Explanation:

B ADSL is correct.

Azure Disk Storage is designed to be used with Azure Virtual Machines, Azure Disk Storage offers high-performance, highly durable block storage for your mission- and business-critical applications. Confidently

migrate to Azure Infrastructure with four disk storage options for the cloud — Ultra Disk Storage, Premium SSD, Standard SSD, and Standard HDD — to optimize costs and performance for your workload needs. Achieve high performance with sub-millisecond latency for throughput and transaction-intensive workloads such as SAP HANA, SQL Server, and Oracle.

<https://azure.microsoft.com/en-in/services/storage/disks/>

Reference:

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

Question: 112

CertyIQ

You need to store data in Azure Blob storage for seven years to meet your company's compliance requirements. The retrieval time of the data is unimportant. The solution must minimize storage costs. Which storage tier should you use?

- A. Archive
- B. Hot
- C. Cool

Answer: A

Explanation:

Archive it to save money

Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements (on the order of hours).

Reference:

https://cloud.netapp.com/blog/azure-blob-storage-pricing-the-complete-guide-azure-cvo-blg#H1_4

Question: 113

CertyIQ

Which type of non-relational data store supports a flexible schema, stores data as JSON files, and stores the all the data for an entity in the same document?

- A. document
- B. columnar
- C. graph
- D. time series

Answer: A

Explanation:

Correct answer is A:document.

Question: 114**CertyIQ**

DRAG DROP -

Match the Azure Cosmos DB APIs to the appropriate data structures.

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

NOTE: Each correct match is worth one point.

Select and Place:

APIs

Cassandra API

Gremlin API

MongoDB API

Table API

Answer Area

Graph data

JSON documents

Key/value data

Answer:**APIs**

Cassandra API

Gremlin API

MongoDB API

Table API

Answer Area

Gremlin API

Graph data

MongoDB API

JSON documents

Table API

Key/value data

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/faq>**Question: 115****CertyIQ**

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:

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).

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.

One advantage of hierarchical namespace is atomic directory manipulation: Object stores approximate a directory hierarchy by adopting a convention of embedding slashes (/) in the object name to denote path segments.

Reference:

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

Question: 116

CertyIQ

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:

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.

Explanation:

As nodes edges by using the Gremlin language.

Reference:

<https://learn.microsoft.com/en-us/azure/cosmos-db/gremlin/introduction>

Question: 117

CertyIQ

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

Answer: B

Explanation:

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.

Enabling redundancy across availability zones makes Azure Cosmos DB resilient to zone outages at the cost of increased charges.

<https://learn.microsoft.com/en-us/azure/cosmos-db/high-availability#replica-outages>

Reference:

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

Question: 118

CertyIQ

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

Answer: D

Explanation:

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: 119

CertyIQ

Your company needs to design a database that shows how changes in network traffic in one area of a network affect network traffic in other areas of the network.

Which type of data store should you use?

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

Answer: A

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: 120

CertyIQ

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 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:

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"/>

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: 121

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	

Answer:

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

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: 122

HOTSPOT -

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

Hot Area:

Answer Area

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

account
container
database
item

Answer:

Answer Area

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

account
container
database
item

Explanation:

Reference:

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

Question: 123

CertyIQ

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

Answer: B

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: 124

CertyIQ

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

Answer: D

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: 125

CertyIQ

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:

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"/>

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: 126

CertyIQ

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:

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).

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-namespace>

Question: 127

CertyIQ

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

Answer: A

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: 128

CertyIQ

You manage an application that stores data in a shared folder on a Windows server. You need to move the shared folder to Azure Storage.

Which type of Azure Storage should you use?

- A. queue

- B. blob
- C. file
- D. table

Answer: C

Explanation:

Azure file shares can be mounted concurrently by cloud or on-premises deployments of Windows, Linux, and macOS. Azure file shares can also be cached on Windows Servers with Azure File Sync for fast access near where the data is being used.

Reference:

<https://azure.microsoft.com/en-us/services/storage/files/>

Question: 129

CertyIQ

Your company is designing a database that will contain session data for a website. The data will include notifications, personalization attributes, and products that are added to a shopping cart. Which type of data store will provide the lowest latency to retrieve the data?

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

Answer: C

Explanation:

Column-family database is the low latency store type.

Examples

- Recommendations
- Personalization
- Sensor data
- Telemetry
- Messaging
- Social media analytics
- Web analytics
- Activity monitoring
- Weather and other time-series data

Question: 130

CertyIQ

You have an application that runs on Windows and requires access to a mapped drive. Which Azure service should you use?

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

Answer: A

Explanation:

Azure Files is Microsoft's easy-to-use cloud file system. Azure file shares can be seamlessly used in Windows and Windows Server.

To use an Azure file share with Windows, you must either mount it, which means assigning it a drive letter or mount point path, or access it via its UNC path.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/files/storage-how-to-use-files-windows>

Question: 131

CertyIQ

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:

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"/>

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: 132

CertyIQ

Your company is designing an application that will write a high volume of JSON data and will have an application-defined schema.

Which type of data store should you use?

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

Answer: C

Explanation:

Key/Value databases can not query json objects, you need to fully retrieve and then process...that means you have really little query options, just a simple group by query will take longer than on a Document type database(the right answer)

A document store such as MongoDB or Couchbase would be a good choice for storing high volumes of JSON data and an application-defined schema.

Document stores are a type of NoSQL database that store data in the form of semi-structured or unstructured documents, which can be represented in JSON format. They allow for storing and querying data in a flexible, schema-less way, which is well-suited for storing high volumes of JSON data with an application-defined schema. They provide built-in support for indexing and querying JSON data, which can be useful for performing complex queries and aggregations on the data.

Alternatively, a key-value store such as Redis or Riak, could also be a good option for storing high volumes of JSON data, as they provide very fast read and write access and are optimized for retrieving data based on a single key. They can be used to store data in memory, which can provide very low latency.

Question: 133

CertyIQ

You need to recommend a non-relational data store that is optimized for storing and retrieving text files, videos, audio streams, and virtual disk images. The data store must store data, some metadata, and a unique ID for each file.

Which type of data store should you recommend?

- A. key/value
- B. columnar
- C. object
- D. document

Answer: C

Explanation:

Object storage is optimized for storing and retrieving large binary objects (images, files, video and audio streams, large application data objects and documents, virtual machine disk images). Large data files are also popularly used in this model, for example, delimiter file (CSV), parquet, and ORC. Object stores can manage extremely large amounts of unstructured data.

Reference:

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

Question: 134

CertyIQ

Your company is designing a data store for internet-connected temperature sensors. The collected data will be used to analyze temperature trends. Which type of data store should you use?

- A. relational
- B. time series
- C. graph
- D. columnar

Answer: B

Explanation:

Time series data is a set of values organized by time. Time series databases typically collect large amounts of data in real time from a large number of sources.

Updates are rare, and deletes are often done as bulk operations. Although the records written to a time-series database are generally small, there are often a large number of records, and total data size can grow rapidly.

Reference:

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

Question: 135

CertyIQ

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
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:

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"/>

Explanation:

Data transfer prices for ADLS

When you write data into GRS accounts, that data will be replicated to another Azure region. The Geo-Replication Data Transfer charge is for the bandwidth of replicating that data to another Azure region. This charge also applies when you change the account replication setting from LRS to GRS or RA-GRS. View the Data transfer prices on Blobs pricing page.

<https://azure.microsoft.com/en-us/pricing/details/storage/data-lake/#pricing>

Azure Blob Storage pricing

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

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: 136

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:**Answer Area**

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

	▼
graph.	
table.	
partition key.	
document.	

Explanation:**API - Gremlin // Container is projected as -Graph // Item is projected as -Nodes and Edges**

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/create-graph-gremlin-console><https://azure.microsoft.com/en-us/blog/a-technical-overview-of-azure-cosmos-db/>**Question: 137**

At which two levels can you set the throughput for an Azure Cosmos DB account? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. database
- B. item

- C. container
- D. partition

Answer: AC

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/set-throughput>

Question: 138

CertyIQ

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 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:

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"/>

Explanation:

Reference:

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

Question: 139

CertyIQ

What is a characteristic of non-relational data?

- A. no indexes
- B. self-describing entities
- C. a well-defined schema
- D. no unique key values

Answer: B

Explanation:

It suggests unique key (which some non-relational stores (like document and key-value) have; relational data has primary and foreign key.

<https://www.coursera.org/lecture/explore-core-data-concepts-microsoft-azure/explore-the-characteristics-of-non-relational-data-nP41V>

Question: 140

CertyIQ

HOTSPOT -

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

Hot Area:

Answer Area

	▼
A hierarchical namespace	
Azure Synapse Link	
PolyBase	
The Data Movement Service (DMS)	

is an Azure Synapse Analytics feature that retrieves data from Azure Blob storage by using Transact-SQL queries.

Answer:

Answer Area

	▼
A hierarchical namespace	
Azure Synapse Link	
PolyBase	
The Data Movement Service (DMS)	

is an Azure Synapse Analytics feature that retrieves data from Azure Blob storage by using Transact-SQL queries.

Explanation:

Polybase

Data virtualization using PolyBase feature is available in preview for Azure SQL Managed Instance, scoped to querying external data stored in files in Azure Data Lake Storage (ADLS) Gen2 and Azure Blob Storage.

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

Reference:

Question: 141

CertyIQ

You need to gather real-time telemetry data from a mobile application. Which type of workload describes this scenario?

- A. Online Transaction Processing (OLTP)
- B. batch
- C. massively parallel processing (MPP)
- D. streaming

Answer: D

Explanation:

"D" as real-time is always Streaming.. OLTP is delayed upto 30 sec

Reference:

<https://docs.microsoft.com/en-in/azure/azure-monitor/overview>

Question: 142

CertyIQ

You have a dedicated SQL pool in Azure Synapse Analytics that is only used actively every night for eight hours. You need to minimize the cost of the dedicated SQL pool as much as possible during idle times. The solution must ensure that the data remains intact.

What should you do on the dedicated SQL pool?

- A. Scale down the data warehouse units (DWUs).
- B. Pause the pool.
- C. Create a user-defined restore point.
- D. Delete the pool

Answer: B

Explanation:

Dedicated SQL pool

You can control costs for a dedicated SQL pool by pausing the resource when it is not in use. For example, if you won't be using the database during the night and on weekends, you can pause it during those times, and resume it during the day. For more information, see [Pause and resume compute in dedicated SQL pool via the Azure portal](#).

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-manage-compute-overview>

Question: 143**CertyIQ**

Which Azure Data Factory component initiates the execution of a pipeline?

- A. a control flow
- B. a trigger
- C. a parameter
- D. an activity

Answer: B**Explanation:**

Trigger.

Pipeline runs are typically instantiated by passing arguments to parameters that you define in the pipeline. You can execute a pipeline either manually or by using a trigger

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/concepts-pipeline-execution-triggers#trigger-execution>

Question: 144**CertyIQ**

Your company has a reporting solution that has paginated reports. The reports query a dimensional model in a data warehouse.

Which type of processing does the reporting solution use?

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

Answer: C**Explanation:**

DWH --> OLAP

Reference:

<https://datawarehouseinfo.com/how-does-oltp-differ-from-olap-database/>

Question: 145**CertyIQ**

DRAG DROP -

Match the types of activities to the appropriate Azure Data Factory activities.

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

NOTE: Each correct match is worth one point.

Select and Place:

Activity Types

Answer Area

Control

Copy

Data movement

Mapping data flow

Data transformation

Lookup

Answer:

Activity Types

Answer Area

Control

Data movement

Copy

Data movement

Data transformation

Mapping data flow

Data transformation

Control

Lookup

Explanation:

Box 1: Data movement -

Box 2: Data transformation -

A pipeline could contain a set of activities that ingest and clean log data, and then kick off a mapping data flow to analyze the log data.

Box 3: Control -

Lookup Activity is a control flow activity.

Lookup Activity can be used to read or look up a record/ table name/ value from any external source. This output can further be referenced by succeeding activities.

Reference:

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

Question: 146

CertyIQ

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

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 on write
- E. schema on read
- F. normalized data

Answer: BDF

Explanation:

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

D: Typical traits of transactional data include: schema on write, strongly enforced

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: 147

CertyIQ

Which two activities can be performed entirely by using the Microsoft Power BI service without relying on Power BI Desktop? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. report and dashboard creation
- B. report sharing and distribution
- C. data modeling
- D. data acquisition and preparation

Answer: AB

Explanation:

Sharing is only done in Power BI Service.

You cannot create a Dashboard in Power BI Desktop as it's created from already existing Reports.

Question: 148

CertyIQ

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:

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.

Explanation:

Control flow (<https://docs.microsoft.com/en-us/azure/data-factory/introduction>)

Control flow is an orchestration of pipeline activities that includes chaining activities in a sequence, branching, defining parameters at the pipeline level, and passing arguments while invoking the pipeline on-demand or from a trigger. It also includes custom-state passing and looping containers, that is, For-each iterators.

Reference:

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

Question: 149

CertyIQ

You need to develop a solution to provide data to executives. The solution must provide an interactive graphical interface, depict various key performance indicators, and support data exploration by using drill down. What should you use in Microsoft Power BI?

- A. a view
- B. a report
- C. a dataflow
- D. Microsoft Power Apps

Answer: B

Explanation:

Reference:

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

Question: 150

CertyIQ

Which two Azure services can be used to provision Apache Spark clusters? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Azure Time Series Insights
- B. Azure HDInsight
- C. Azure Databricks
- D. Azure Log Analytics

Answer: BC

Explanation:

Reference:

<https://www.sqlshack.com/a-beginners-guide-to-azure-databricks/>

Question: 151

CertyIQ

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

Answer: C

Explanation:

Quality assurance application focuses on improving the software development process and making it efficient and effective as per quality standards.

OLAP applies complex queries to large amounts of historical data, aggregated from OLTP databases and other sources, for data mining, analytics, and business intelligence projects.

Moreover, the purpose of OLAP ----> Plan, solve problems, support decisions, discover hidden insights.

According to the question, the application is reading data from the data warehouse for quality assurance which is clearly - discovering insights from the data. Thus, the answer is OLAP.

Question: 152

CertyIQ

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

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

Answer: ACE

Explanation:

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: 153

CertyIQ

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:

Answer Area

Statements	Yes	No
A Microsoft Power BI dashboard is associated with a single workspace.	<input checked="" type="radio"/>	<input 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"/>

Explanation:

ONE dashboard...

- is associated with a single workspace
- can display visualizations from many different datasets
- can display visualizations from many different reports
- can display visualizations pinned from other tools (for example, Excel)

Source:

<https://docs.microsoft.com/en-us/power-bi/fundamentals/service-basic-concepts#dashboards>

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: 154

CertyIQ

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

- A. SSIS packages
- B. an integration runtime
- C. a control flow
- D. a pipeline

Answer: B

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: 155

CertyIQ

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:

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.

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: 156

CertyIQ

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. Enforce business logic across reports

Answer: AC

Explanation:

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: 157

CertyIQ

You need to use Transact-SQL to query files in Azure Data Lake Storage Gen 2 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

Answer: C

Explanation:

Polybase is the correct answer,

<https://docs.microsoft.com/en-gb/learn/modules/explore-data-storage-processing-azure/3-explore-azure-synapse-analytics>

"Use SQL pools in Synapse Analytics for the following scenarios:

Complex reporting. You can use the full power of Transact-SQL to run complex SQL statements that summarize and aggregate data.

Data ingestion. PolyBase enables you to retrieve data from many external sources and convert it into a tabular format. You can reformat this data and save it as tables and materialized views in Azure Synapse."

Reference:

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

Question: 158

CertyIQ

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 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 BI dashboard can display visualizations from a Microsoft Excel workbook.	<input type="radio"/>	<input type="radio"/>

Answer:

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 BI dashboard can display visualizations from a Microsoft Excel workbook.	<input checked="" type="radio"/>	<input type="radio"/>

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: 159

CertyIQ

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

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

Answer: BDF

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: 160

CertyIQ

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

Answer: A

Explanation:

The primary purpose of a data warehouse is to provide answers to complex queries that rely on data from multiple sources.

A data warehouse is a large, centralized repository of data that is specifically designed for reporting and analysis. The primary purpose of a data warehouse is to store historical data from multiple sources in a single location, in a format that makes it easy to query and analyze. This data is typically read-only, meaning that it is not intended for transactional use. Instead, it is used for reporting and analysis, such as generating business intelligence reports, performing data mining and modeling, and supporting decision-making processes. Data warehousing also includes transforming data from various sources into a consistent format, making it easy for users to query and analyze data from multiple systems.

Question: 161

CertyIQ

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 Cognitive Search

Azure Data Catalog

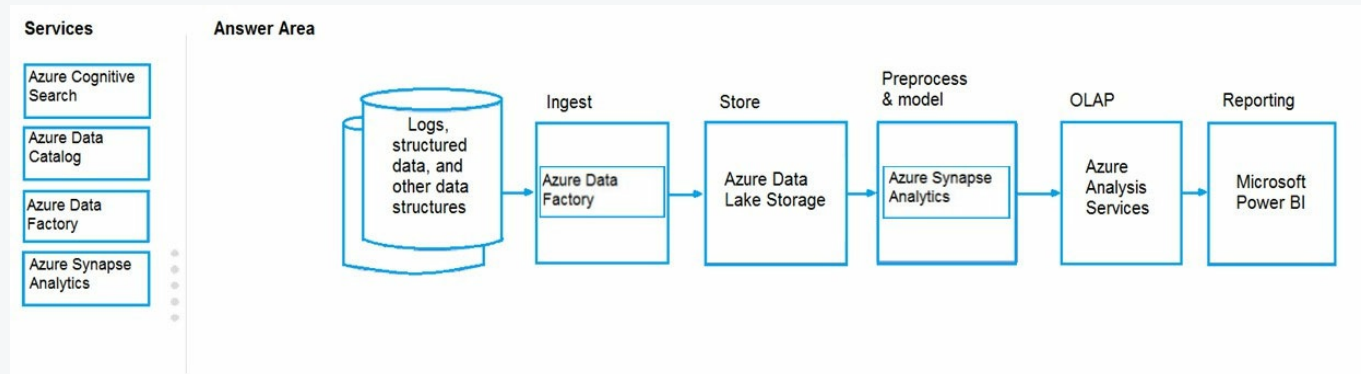
Azure Data Factory

Azure Synapse Analytics

Answer Area



Answer:



Explanation:

Box Ingest: Azure Data Factory -

You can build a data ingestion pipeline with Azure Data Factory (ADF).

Box Preprocess & model: Azure Synapse Analytics

Use Azure Synapse Analytics to preprocess data and deploy machine learning models.

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/how-to-data-ingest-adf> <https://docs.microsoft.com/en-us/azure/machine-learning/team-data-science-process/sqldw-walkthrough>

Question: 162

CertyIQ

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
- Streaming

- Workload type Analyzing historical data containing web traffic collected during the previous year.
- Workload type Classifying images that were uploaded last month.
- Workload type Tracking in real time how many people are currently using a website.

Answer:

Workload Types

Answer Area

- Batch
- Streaming

- Batch Analyzing historical data containing web traffic collected during the previous year.
- Batch Classifying images that were uploaded last month.
- Streaming Tracking in real time how many people are currently using a website.

Explanation:

Box 1: Batch -

The batch processing model requires a set of data that is collected over time while the stream processing model requires data to be fed into an analytics tool, often in micro-batches, and in real-time. The batch Processing model handles a large batch of data while the Stream processing model handles individual records or micro-batches of few records. In Batch Processing, it processes over all or most of the data but in Stream Processing, it processes over data on a rolling window or most recent record.

Box 2: Batch -

Box 3: Streaming -

Reference:

<https://k21academy.com/microsoft-azure/dp-200/batch-processing-vs-stream-processing>

Question: 163

CertyIQ

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:

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

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: 164

CertyIQ

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 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:

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"/>

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: 165

CertyIQ

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 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

Answer:

Services

Azure Data Lake Storage

Azure SQL Database

Azure Synapse Analytics

Answer Area

Azure Synapse Analytics

Output data to Parquet format

Azure Data Lake Storage

Store data that is in Parquet format

Azure SQL Database

Persist a tabular representation of data that is stored in Parquet format

Explanation:

- Azure Synapse Analytics: Output data to parquet format
- Azure Data Lake Storage: Store data that is in parquet format
- Azure SQL Database: Persist a tabular representation of data that is stored in parquet format

Reference:

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

Question: 166

CertyIQ

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 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:

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"/>

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: 167

CertyIQ

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

Answer: D

Explanation:

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: 168

CertyIQ

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

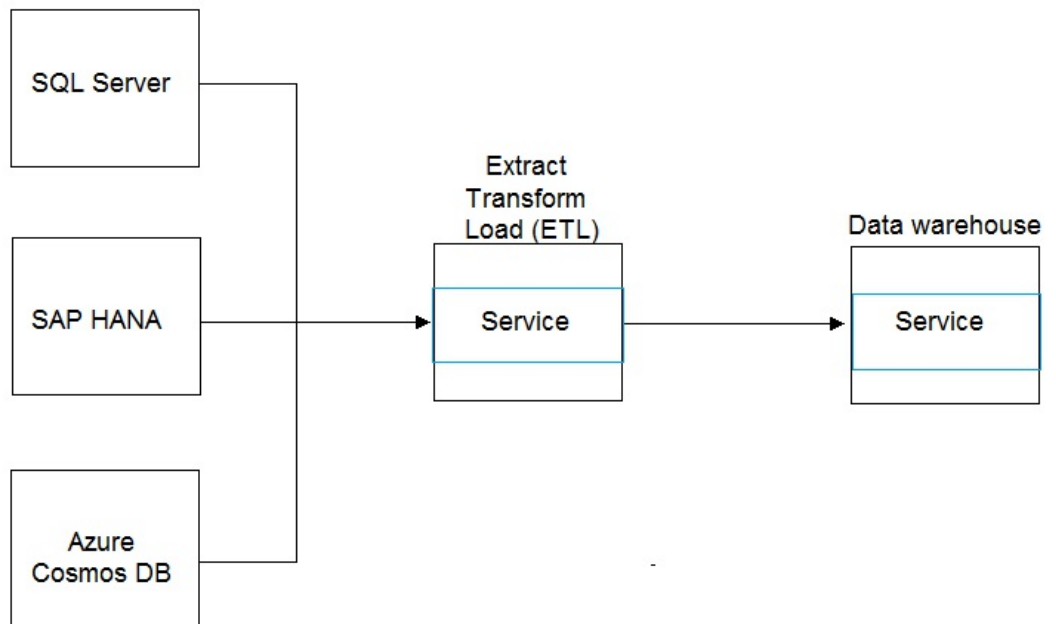
Azure Data Factory

Azure Table storage

Azure Cosmos DB

Azure Synapse Analytics

Answer Area



Answer:

Services

Azure Analysis Services

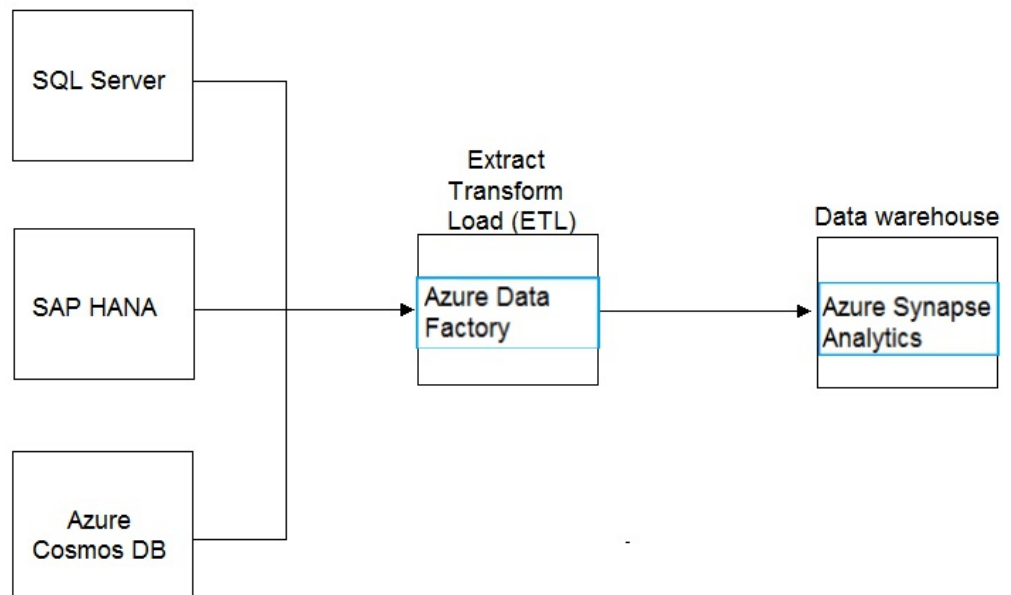
Azure Data Factory

Azure Table storage

Azure Cosmos DB

Azure Synapse Analytics

Answer Area



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: 169

CertyIQ

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:

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"/>

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: 170

CertyIQ

Which scenario is an example of a streaming workload?

- A. sending transactions that are older than a month to an archive

- B. sending transactions daily from point of sale (POS) devices
- C. sending telemetry data from edge devices
- D. sending cloud infrastructure metadata every 30 minutes

Answer: C

Explanation:

Telemetry is the in situ collection of measurements or other data at remote points and their automatic transmission to receiving equipment (telecommunication) for (real-time) monitoring.

Question: 171

CertyIQ

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:

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.

Explanation:

Option A is wrong why? It's simple, batch workloads always collect all the data first before processing

Option B is wrong. Although its collecting and processing data once a day it's not the best answer, it's obvious that D will be the best answer for this question

Option C is also wrong. It's real time so batch workloads won't work that way

Option D is the correct answer it has the same explanation with Option B but option B only tells that batch workloads only happens a day well that's not always the case since batch workloads may vary on many things sometimes it can process the data after 12 hours, 1 week, 1 month or a certain amount of data workload which explains why there is a condition

Question: 172

CertyIQ

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
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:

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"/>

Explanation:

Yes | Yes | No

C is No because the electric meter does not send EVERY change it registers immediately, it just gives you daily data, which might contain just one reading, or hourly readings, or readings for each minute.

Question: 173

CertyIQ

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

Answer: C

Explanation:

Measures of Central Tendency (Mean, Median, and Mode).

Measures of Dispersion or Variation (Variance, Standard Deviation, Range).

<https://www.intellspot.com/descriptive-statistics-examples/>

Question: 174

CertyIQ

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:

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 checked="" type="radio"/>	<input type="radio"/>

Explanation:

1. Stream processing has access to the most recent data received or data within a rolling time window.

Answer: Yes

- Stream processing continuously processes data as it arrives, often within a rolling time window. This enables real-time or near-real-time analysis of the most recent data.

2. Batch processing must occur immediately and have latency in the order of seconds or milliseconds.

Answer: No

- Batch processing processes large volumes of data at once and is not time-sensitive. It typically has higher latency, ranging from minutes to hours, and does not need to occur immediately.

3. Stream processing is used for simple response functions, aggregates, or calculations such as rolling averages.

Answer: Yes

- Stream processing is ideal for real-time calculations like rolling averages, aggregates, or event-based response functions, which are performed on data streams as they flow in.

Key Tip:

Stream processing is suited for real-time, low-latency tasks on continuous data streams.

Batch processing is suited for high-volume, non-urgent computations with longer latency.

Reference:

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

Question: 175

CertyIQ

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

Answer: C

Explanation:

Synapse Link is a feature of Azure Cosmos DB that allows you to enable real-time analytics on your Cosmos

DB data, by creating a seamless connection between Azure Cosmos DB and Azure Synapse Analytics. It enables you to run analytical queries against your Cosmos DB data using Synapse SQL pool, and perform complex joins and aggregations across data stored in both Cosmos DB and other data sources.

Reference:


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

Question: 176

CertyIQ

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

Sales by Year and Quarter



What should you create in Power BI Desktop?

- A. a waterfall chart
- B. a ribbon chart
- C. a bar chart
- D. a decomposition tree

Answer: A

Explanation:

Reference:

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

Question: 177

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

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:

Answer Area

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

	▼
graph	
table	
partition key	
document	

Question: 178

CertyIQ

You have the following two tables of customer data.

CustomerID	NameStyle	Title	FirstName	MiddleName	LastName
29486	0	Ms.	Kim	NULL	Abercrombie
29489	0	Ms.	Frances	B.	Adams
29490	0	Ms.	Margaret	J.	Smith
29492	0	Mr.	Jay	NULL	Adams

CustomerID	AddressID	AddressType	ModifiedDate
29486	621	Main Office	9/1/05 0:00
29489	1069	Main Office	7/1/05 0:00
29490	887	Main Office	9/1/06 0:00
29492	618	Main Office	12/1/06 0:00
29494	537	Main Office	9/1/05 0:00

Which type of data do the tables contain?

- A. structured
- B. non-relational
- C. semi-structured
- D. unstructured

Answer: A

Explanation:

it is Structured. it have Primary and Foreign Key Constraints

<https://learn.microsoft.com/en-us/sql/relational-databases/tables/primary-and-foreign-key-constraints?view=sql-server-ver16>

Question: 179

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

You can query a graph database in Azure Cosmos DB

	▼
as nodes and edges by using the Gremlin language	
as partitioned row store by using Cassandra Query Language (CQL)	
as partitioned row store by using Language-Integrated Query (LINQ)	

Answer:

Answer Area

You can query a graph database in Azure Cosmos DB

- as nodes and edges by using the Gremlin language
- as partitioned row store by using Cassandra Query Language (CQL)
- as partitioned row store by using Language-Integrated Query (LINQ)

Explanation:

The Gremlin API in Azure Cosmos DB uses the Gremlin query language.

<https://learn.microsoft.com/en-us/azure/cosmos-db/gremlin/introduction>

Question: 180

CertyIQ

Which Azure Blob Storage feature enables you to sync a Blob Storage account to an external cache?

- A. the Hot access tier
- B. change feed
- C. snapshots
- D. versioning

Answer: B

Explanation:

Change feed

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed>

Change feed support is well-suited for scenarios that process data based on objects that have changed. For example, applications can:

- Update a secondary index, synchronize with a cache, search-engine, or any other content-management scenarios.

Question: 181

CertyIQ

You have an on-premises Microsoft SQL Server database.

You need to migrate the database to the cloud. The solution must meet the following requirements:

- * Minimize maintenance effort.
- * Support the Database Mail and Service Broker features.

What should you include in the solution?

- A. Azure SQL Database single database
- B. an Azure SQL Database elastic pool
- C. Azure SQL Managed instance
- D. SQL Server on Azure virtual machines

Answer: C

Explanation:

Azure SQL Database does not support service broker and mail. You need the Managed Instance.

<https://docs.microsoft.com/en-us/azure/azure-sql/database/features-comparison?view=azuresql>

Question: 182

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

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

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

Answer:

Answer Area

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

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

Question: 183

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

Batch workloads

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

Answer:

Answer Area

Batch workloads

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

Question: 184

CertyIQ

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

A standalone data analysis tool

The CRM system

The data warehouse

Answer Area

Extract:

Location

Load:

Location

Transform:

Location

Answer:

Locations

A standalone data analysis tool

The CRM system

The data warehouse

Answer Area

Extract:

The CRM system

Load:

The data warehouse

Transform:

The data warehouse

Explanation:

Extract: The CRM system

In the ETL (Extract, Transform, Load) process, data is first extracted from the source system. In this case, the CRM system is the source where customer data is stored, making it the logical choice for extraction.

Load: The data warehouse

After transformation, the data is loaded into the target system for storage and analysis. The data warehouse is the destination where the transformed data is stored for further querying and reporting.

Transform: The data warehouse

While transformation can occur in various stages, it often happens either in an intermediary step or directly within the data warehouse to format, cleanse, or aggregate data to meet analytical requirements.

Key Tip:

ETL Process:

Extract: Identify and pull data from source systems like CRM, ERP, or databases.

Transform: Cleanse, format, and aggregate data for analysis, often within the warehouse or during intermediate stages.

Load: Place the prepared data into a target system, typically a data warehouse, for querying and business intelligence tasks.

Question: 185

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

Descriptive analytics tells you

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

Answer:

Answer Area

Descriptive analytics tells you

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

Question: 186

CertyIQ

What is a function of a modern data warehouse?

- A. supports batch processing only
- B. supports real-time and batch processing
- C. provides built-in or native online analytical processing
- D. stores raw data only

Answer: B

Explanation:

Personal preference goes to B. Key phrase is 'modern' data warehouse.

From Microsoft learn:

Large-scale data warehousing solutions combine conventional data warehousing used to support business intelligence (BI) with techniques used for so-called "big data" analytics. A conventional data warehousing solution typically involves copying data from transactional data stores into a relational database with a schema that's optimized for querying and building multidimensional models. Big Data processing solutions on the other hand, are used with large volumes of data in multiple formats, which is batch loaded or captured in real-time streams and stored in a data lake from which distributed processing engines like Apache Spark are used to process it.

Answer C is functional of conventional data warehouse.

Reference

<https://learn.microsoft.com/bs-latn-ba/training/modules/examine-components-of-modern-data-warehouse/1-introduction>

Question: 187

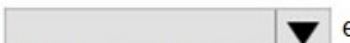
HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer AreaAn Azure Storage account must have  enabled to support Azure Data Lake Storage.

- Blob soft delete
- Hierarchical namespace
- Large file shares
- Versioning

Answer:**Answer Area**An Azure Storage account must have  enabled to support Azure Data Lake Storage.

- Blob soft delete
- Hierarchical namespace
- Large file shares
- Versioning

Explanation:

hierarchical namespace

<https://docs.microsoft.com/en-us/azure/storage/blobs/create-data-lake-storage-account>**Question: 188**

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

Answer: A**Explanation:**

Notebooks are a common tool in data science and machine learning for developing code and presenting results. In Azure Databricks, notebooks are the primary tool for creating data science and machine learning workflows and collaborating with colleagues. Databricks notebooks provide real-time coauthoring in multiple languages, automatic versioning, and built-in data visualizations.

Reference

Introduction to Databricks notebooks

<https://learn.microsoft.com/en-us/azure/databricks/notebooks/>

Question: 189

CertyIQ

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. HTML
- C. YAML
- D. JSON

Answer: D**Question: 190**

CertyIQ

You have data saved in the following format.

```
<Person Age= "23">
  <FirstName>John</FirstName>
  <LastName>Smith</LastName>
  <Hobbies>
    <Hobby Type= "Sports">Golf</Hobby>
    <Hobby Type= "Leisure">Reading</Hobby>
    <Hobby Type= "Leisure">Guitar</Hobby>
  </Hobbies>
</Person>
```

Which format was used?

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

Answer: D

Question: 191

CertyIQ

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

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

Answer: C

Question: 192

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

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:

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.	

Explanation:

A relational database should be used when strong consistency guarantees are required and when a dynamic schema is required.

Relational databases are designed to store data in a structured way, using tables, rows, and columns. They have a fixed schema that defines the structure of the data, which makes them well-suited for situations where data needs to be stored in a structured and predictable way, and where consistency of the data is important. They also allow to enforce constraints and relationships between the data, which can provide additional guarantees of data integrity.

Question: 193

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

Relational data is stored in

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

Answer:

Answer Area

Relational data is stored in

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

Question: 194

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

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

Answer: A

Explanation:

Atomicity is the right answer.

An atomic transaction is an indivisible and irreducible series of database operations such that either all occurs, or nothing occurs. A guarantee of atomicity prevents updates to the database occurring only partially,

which can cause greater problems than rejecting the whole series outright. As a consequence, the transaction cannot be observed to be in progress by another database client.

Question: 195

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

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

is an object associated with a table that sorts on their key values.

Answer:

Answer Area

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

is an object associated with a table that sorts on their key values.

Question: 196

CertyIQ

Which Azure Storage service implements the key/value model?

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

Answer: C

Explanation:

C - <https://learn.microsoft.com/en-us/azure/storage/tables/table-storage-overview>

Azure Table storage is a service that stores non-relational structured data (also known as structured NoSQL data) in the cloud, providing a key/attribute store with a schemaless design.

Question: 197

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

Social media users and followers in a graph database is an example of

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

Answer:

Answer Area

Social media users and followers in a graph database is an example of

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

Explanation:

Semi-structured data is correct answer.

Question: 198

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

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:

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.	

Question: 199

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

In batch processing,

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

Answer:

Answer Area

In batch processing,

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

Explanation:

Batch processing involves collecting data over a period of time and processing it as a group (batch). This inherently introduces delays because the processing does not occur in real-time; instead, it waits for the batch to be complete.

Why the other answers are incorrect:

"Data is always inserted one row at a time":

Batch processing typically involves handling data in bulk, not row-by-row insertion. This is more characteristic of transactional or real-time systems.

"Data is processed in real-time":

Real-time processing is the opposite of batch processing. In batch processing, there is no immediate processing as the focus is on efficiency over timeliness.

"Processing can only execute serially":

Batch processing can involve parallelism depending on the system design, so this is not an absolute characteristic.

Important Tip:

Batch processing is ideal for scenarios where immediate results are not required, but efficiency in handling large data volumes is crucial (e.g., payroll processing, billing systems).

Question: 200**CertyIQ**

You plan to deploy an app. The app requires a nonrelational data service that will provide latency guarantees of less than 10-ms for reads and writes.

What should you include in the solution?

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

Answer: D**Explanation:****D. Azure Cosmos DB**

The read latency for all consistency levels is always guaranteed to be less than 10 milliseconds at the 99th percentile. The average read latency, at the 50th percentile, is typically 4 milliseconds or less.

The write latency for all consistency levels is always guaranteed to be less than 10 milliseconds at the 99th

percentile. The average write latency, at the 50th percentile, is usually 5 milliseconds or less. Azure Cosmos DB accounts that span several regions and are configured with strong consistency are an exception to this guarantee.

<https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels>

Question: 201

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A non-relational database

- imposes a fixed schema on data
- store entities without a unique key value
- stores data in relational tables
- stores data without a fixed schema

Answer:

Answer Area

A non-relational database

- imposes a fixed schema on data
- store entities without a unique key value
- stores data in relational tables
- stores data without a fixed schema

Explanation:

Answer is D.

Non relational DATA had no unique key value.

Non-relational DATABASE stores data without a fixed schema

Question: 202

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:

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 checked="" 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 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"/>

Explanation:

Yes

No

Yes

Reference:

<https://learn.microsoft.com/en-us/training/modules/explore-core-data-concepts/5-transactional-data-processing>

Question: 203

You have data saved in the following format.

```
FirstName, LastName, Age, LeisureHobby, SportsHobby
John, Smith, 23, Reading, Backteball
Ben, Smith, 21, Guitar Curling
```

Which format was used?

- A. YAML
- B. CSV
- C. JSON
- D. HTML

Answer: B

Explanation:

CSV

YAML Syntax:

nearly every YAML file starts with a list. Each item in the list is a list of key/value pairs, commonly called a “hash” or a “dictionary”. So, we need to know how to write lists and dictionaries in YAML.

There’s another small quirk to YAML. All YAML files (regardless of their association with Ansible or not) can optionally begin with --- and end with This is part of the YAML format and indicates the start and end of a document.

All members of a list are lines beginning at the same indentation level starting with a "- " (a dash and a space):

A list of tasty fruits

- Apple
- Orange
- Strawberry
- Mango

...

A dictionary is represented in a simple key: value form (the colon must be followed by a space):

An employee record

martin:

name: Martin D'vloper

job: Developer

skill: Elite

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

Answer Area

▼

 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:

Answer Area

▼

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

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

Question: 205

CertyIQ

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

Answer Area

Backing up and restoring data are the responsibilities of a

▼

data analyst
data engineer
data scientist
database administrator

Answer:

Answer Area

Backing up and restoring data are the responsibilities of a

▼

data analyst
data engineer
data scientist
database administrator

Explanation:

it is Database administrators.

Database administrators manage databases, assigning permissions to users, storing backup copies of data and restore data in the event of a failure.

<https://learn.microsoft.com/en-us/training/modules/explore-roles-responsibilities-world-of-data/2-explore-job-roles>

Question: 206

CertyIQ

What is required to provision Azure Data Lake Storage in an Azure Storage account?

- A. Versioning must be disabled.
- B. Hierarchical namespace must be disabled.
- C. Versioning must be enabled.
- D. Hierarchical namespace must be enabled.

Answer: D

Explanation:

Reference:

<https://azure.microsoft.com/en-us/updates/general-availability-enable-hierarchical-namespace-for-existing-azure-storage-accounts/>

Question: 207

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

Microsoft SQL Server database tables store

	▼
non-relational data.	
relational data.	
semi-structured data.	

Answer:

Answer Area

Microsoft SQL Server database tables store

	▼
non-relational data.	
relational data.	
semi-structured data.	

Explanation:

Microsoft SQL Server database tables store: relational data.

Question: 208

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

	▼
An Azure SQL Database elastic pool	
An instance of SQL Server on Azure Virtual Machines	
A single database in Azure SQL Database	

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

Answer:

Answer Area

	▼
An Azure SQL Database elastic pool	
An instance of SQL Server on Azure Virtual Machines	
A single database in Azure SQL Database	

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

Explanation:

An instance of SQL Server on Azure Virtual Machines.

Question: 209

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

In Microsoft Power BI, a

	▼
dashboard	
dataset	
report	
tile	

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

Answer:

Answer Area

In Microsoft Power BI, a

	▼
dashboard	
dataset	
report	
tile	

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

dashboard
dataset
report
tile

Explanation:

C: Reports

****Report****, the question mentioned more than "one page", and dashboard only use one page.

<https://learn.microsoft.com/en-us/power-bi/create-reports/service-dashboards#dashboards-versus-reports>

Look at the link. there is a table where it's written:

Capability/Dashboards/Reports

Pages/One page/One or more pages

Question: 210

CertyIQ

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

Data ingestion

Extract, load, and transform (ELT) processing

Extract, transform, and load (ETL) processing

Answer Area

Data is retrieved and processed before being saved to a target system

Data is saved to a target system, and then processed

Data is extracted from a source system, and then saved to a target system

Answer:

Processes

Data ingestion

Extract, load, and transform (ELT) processing

Extract, transform, and load (ETL) processing

Answer Area

Data is retrieved and processed before being saved to a target system

Extract, transform, and load (ETL) processing

Data is saved to a target system, and then processed

Extract, load, and transform (ELT) processing

Data is extracted from a source system, and then saved to a target system

Data ingestion

Explanation:

ETL

ELT

Data Ingestion.

Question: 211

CertyIQ

What is used to define a query in a stream processing jobs in Azure Stream Analytics?

- A. YAML
- B. KQL
- C. SQL
- D. XML

Answer: C

Explanation:

Correct Answer C

Azure Stream Analytics uses a SQL query language

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-introduction>

Question: 212

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

	▼
An Azure SQL server	
An elastic pool	
The Data Migration Assistant (DMA)	
A virtual network	

enables multiple Azure SQL databases to share resources such as memory, data storage space, and processing power.

Answer:

	▼
An Azure SQL server	
An elastic pool	
The Data Migration Assistant (DMA)	
A virtual network	

enables multiple Azure SQL databases to share resources such as memory, data storage space, and processing power.

Question: 213

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

XML-formatted files are examples of

relational data.
semi-structured data.
unstructured data.

Answer:

XML-formatted files are examples of

relational data.
semi-structured data.
unstructured data.

Explanation:

"Semi-Structured"

<https://learn.microsoft.com/en-us/training/modules/explore-core-data-concepts/2-data-formats>

Semi-structured data is information that has some structure, but which allows for some variation between entity instances. For example, while most customers may have an email address, some might have multiple email addresses, and some might have none at all.

One common format for semi-structured data is JavaScript Object Notation (JSON). The example below shows a pair of JSON documents that represent customer information. Each customer document includes address and contact information, but the specific fields vary between customers.

Question: 214

CertyIQ

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

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

Answer: D

Explanation:

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

Question: 215

CertyIQ

What is a primary characteristic of a relational database?

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

Answer: B

Explanation:

We can Normalized it to remove the duplicate, so the answer left is B. SQL

Question: 216

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

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

Answer:

Answer Area

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

Question: 217

CertyIQ

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

- A. GRANT
- B. INSERT
- C. REVOKE
- D. DROP

Answer: B

Explanation:

INSERT is a right answer.

Question: 218**CertyIQ**

What is a characteristic of a non-relational database?

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

Answer: C**Explanation:**

<https://www.coursera.org/lecture/explore-core-data-concepts-microsoft-azure/explore-the-characteristics-of-non-relational-data-nP41V>

Question: 219**CertyIQ**

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

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

Answer: D**Explanation:**

<https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview>

Question: 220**CertyIQ**

In Azure Table storage, each row in a table must be uniquely identified by which two components? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. a partition key
- B. a range
- C. a row key
- D. a timestamp

Answer: AC**Explanation:**

From: <https://learn.microsoft.com/en-us/azure/storage/tables/table-storage-overview>

Properties: A property is a name-value pair. Each entity can include up to 252 properties to store data. Each entity also has three system properties that specify a partition key, a row key, and a timestamp. Entities with

the same partition key can be queried more quickly, and inserted/updated in atomic operations. An entity's row key is its unique identifier within a partition.

Question: 221

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

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

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

Answer:

Answer Area

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

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

Explanation:

Data analysts: explore and analyze data to create visualizations and charts that enable organizations to make informed decisions.

Data visualization engineer

Data visualization engineers (or specialists) are responsible for the design and development of data visualizations and dashboards using data sets from various sources. They are experts in data and storytelling and whizzes at business intelligence (BI) tools, often creating user-friendly dashboards for business teams. Data visualization engineers may also write and test analytics platform code, toggle with software tools like Tableau or Looker to produce dashboards, and fix bugs whenever necessary.

<https://www.coursera.org/articles/data-visualization-jobs>

Question: 222

CertyIQ

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

Answer: B

Explanation:

B. Azure Queue storage.

While both Azure Table storage and Azure Queue storage can store semi-structured data, Azure Queue

storage is better suited for event log data that is received in real-time. It provides a message-oriented storage model that allows for reliable and scalable delivery of event logs as messages. Azure Queue storage is also easy to use and can be integrated with other Azure services for further processing and analysis.

Question: 223

CertyIQ

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

- A. Table
- B. Cassandra
- C. Core (SQL)
- D. Gremlin

Answer: D

Explanation:

Gremlin is a right answer.

Question: 224

CertyIQ

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
Azure Data Lake storage can be used as a staging area for ingested data before the data is converted into a format suitable for performing analytics.	<input type="radio"/>	<input type="radio"/>
Azure Data Factory is an analytics engine that uses a massively parallel processing (MPP) architecture.	<input type="radio"/>	<input type="radio"/>
Azure HDInsight enables you to build tabular models from multiple data sources to support online analytical processing (OLAP) queries.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements	Yes	No
Azure Data Lake storage can be used as a staging area for ingested data before the data is converted into a format suitable for performing analytics.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Data Factory is an analytics engine that uses a massively parallel processing (MPP) architecture.	<input type="radio"/>	<input checked="" type="radio"/>
Azure HDInsight enables you to build tabular models from multiple data sources to support online analytical processing (OLAP) queries.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:

Y - Azure Data Lake Storage can be used as a staging area for ingested data before the data is converted into a format suitable for performing analytics. It allows you to store and analyze large amounts of data, and it is optimized for big data analytics workloads.

N - Azure Data Factory is a cloud-based data integration service that allows you to create, schedule, and manage data pipelines. It does not use a massively parallel processing architecture by default, but allows you

to perform complex data integration and transformation tasks by integrating with other Azure services.

N - Azure HDInsight is designed for big data analytics workloads and does not enable you to build tabular models from multiple data sources to support OLAP queries. Azure Analysis Services is a service that you can use to create and manage tabular models for OLAP.

Question: 225

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

An Azure SQL server
An elastic pool
The Data Migration Assistant (DMA)
A virtual network

enables multiple Azure SQL databases to share resources such as memory, data storage space, and processing power.

Answer:

Answer Area

An Azure SQL server
An elastic pool
The Data Migration Assistant (DMA)
A virtual network

enables multiple Azure SQL databases to share resources such as memory, data storage space, and processing power.

Question: 226

CertyIQ

A bank has a system that manages financial transactions.

When transferring money between accounts, the system must never retrieve a value for the source account that reflects the balance before the transfer and a value for the destination account that reflects the balance after the transfer.

Of which ACID semantic is this an example?

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

Answer: D

Explanation:

Isolation.

See course material:

o concurrent transactions cannot interfere with one another, and must result in a consistent database state.

o For example, while the transaction to transfer funds from one account to another is in-process, another transaction that checks the balance of these accounts must return consistent results - the balance-checking transaction can't retrieve a value for one account that reflects the balance before the transfer, and a value for the other account that reflects the balance after the transfer.

Question: 227

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.

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)	

Answer:

Extract, transform, and load (ETL)

Fact tables

Data warehouses

Explanation:

ETL

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

Fact table

<https://www.techtarget.com/searchdatamanagement/definition/fact-table#:~:text=A%20fact%20table%20stores%20quantitative,fact%20table%20can%20be%20analyzed.>

DW

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

Question: 228

HOTSPOT

-

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.

account
container
database
item

Answer:

Answer Area

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

account
container
database
item

Explanation:

Account.

Question: 229

CertyIQ

A data engineer is responsible for which task?

- A. Explore data to identify trends.
- B. Implement policies, tools, and processes for backup and recovery plans.
- C. Design and build analytical models.
- D. Design and implement data stores for analytical workloads.

Answer: D

Explanation:

B = DBA

D = Data Engineer

A data engineer collaborates with stakeholders to design and implement data-related workloads, including data ingestion pipelines, cleansing and transformation activities, and data stores for analytical workloads. They use a wide range of data platform technologies, including relational and non-relational databases, file stores, and data streams.

<https://learn.microsoft.com/en-us/training/modules/explore-roles-responsibilities-world-of-data/2-explore-job-roles>

<https://learn.microsoft.com/en-us/training/modules/explore-roles-responsibilities-world-of-data/2-explore-job-roles>

Question: 230

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

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

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

Answer:

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

distributes processing across compute nodes.

Explanation:

Question: 231

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

Data in a data warehouse can be aggregated and loaded into an online analytical processing (OLAP) model,
which is also known as a

- cube.
- data lake.
- dimension.
- measure.

Answer:

is also known as a ▼

Explanation:

Cube

Question: 232

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

Bar charts
Pie charts
Scatter charts
Tables

are used to compare two numeric measures and identify a relationship or correlation between them.

Answer:

 ▼ are used to

Explanation:

Scatter Charts - A scatter chart (or scatter graph) uses dots to represent values for two different numeric variables. The position of each dot on the horizontal and vertical axis indicates values for an individual data point. Scatter plots are used to observe relationships between variables.

Question: 233

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

A block of code that runs in a database is called

- a stored procedure.
- a table.
- a view.
- an index.

Answer:

a stored procedure.

Explanation:

a stored procedure.

Question: 234

CertyIQ

Which file format has defined names and data types for each column and uses compressed columnar storage?

- A. HTML
- B. CSV
- C. Apache Parquet
- D. JSON

Answer: C

Explanation:

Key phrase is "compressed columnar storage". That is about Parquet

Question: 235

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

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:

Explanation:

Transactional writes is a right answer.

Question: 236

CertyIQ

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.

Answer Area

Statements	Yes	No
Data stored in the Hot access tier of Azure Blob storage is stored on high-performance media.	<input type="radio"/>	<input type="radio"/>
The Cool access tier in Azure Blob storage incurs higher storage costs compared to the Hot access tier.	<input type="radio"/>	<input type="radio"/>
The Archive access tier in Azure Blob storage has the highest latency of the three available tiers.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements	Yes	No
Data stored in the Hot access tier of Azure Blob storage is stored on high-performance media.	<input checked="" type="radio"/>	<input type="radio"/>
The Cool access tier in Azure Blob storage incurs higher storage costs compared to the Hot access tier.	<input type="radio"/>	<input checked="" type="radio"/>
The Archive access tier in Azure Blob storage has the highest latency of the three available tiers.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Y/N/Y

<https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview>

The archive tier is an offline tier for storing data that is rarely accessed. The archive access tier has the lowest storage cost. However, this tier has higher data retrieval costs with a higher latency as compared to the hot and cool tiers.

Question: 237

CertyIQ

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.

Answer Area

Statements	Yes	No
Online transaction processing (OLTP) relies on a database system where data storage is optimized for both read and write operations.	<input type="radio"/>	<input type="radio"/>
Online analytical processing (OLAP) relies on a database system where data storage is optimized for both read and write operations.	<input type="radio"/>	<input type="radio"/>
Data warehouses are optimized for read operations.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements	Yes	No
Online transaction processing (OLTP) relies on a database system where data storage is optimized for both read and write operations.	<input checked="" type="radio"/>	<input type="radio"/>
Online analytical processing (OLAP) relies on a database system where data storage is optimized for both read and write operations.	<input type="radio"/>	<input checked="" type="radio"/>
Data warehouses are optimized for read operations.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

The answer be YNY

OLTP is for relational database where we do both type(read/write) of operations.

Question: 238

CertyIQ

A retail point of sale (POS) system is an example of which type of solution?

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

Answer: D

Explanation:

D is correct: got to be OLTP

A retail point of sale (POS) system is an example of an Online Transaction Processing (OLTP) solution.

Question: 239

CertyIQ

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

In a graph database, edges specify relationships between

▼

column families
facts
vertexes
tables

Answer:

Answer Area

In a graph database, edges specify relationships between

▼

column families
facts
vertexes
tables

Explanation:

A graph database is a collection of nodes (or vertices) and edges (or relationships).

Therefore, the correct answer is Vertexes.

Question: 240

CertyIQ

You have structured data in tabular format.

What represents an individual instance of a data entity?

- A.a column
- B.a table
- C.a cell
- D.a row

Answer: D

Explanation:

In a simple relational database implementation, each row of a table represents one instance of an entity type, and each field in a table represents an attribute type.

An entity is a thing in which information needs to be known or held. A table is the object that stores a collection of entities. A row represents a single instance of an entity. A column defines a specific property of an entity. A row.

Question: 241

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

Relational databases are commonly used to store and query

▼

blobs.

semi-structured data.

structured data.

unstructured data.

Answer:

Answer Area

Relational databases are commonly used to store and query

▼

blobs.

semi-structured data.

structured data.

unstructured data.

Question: 242

CertyIQ

Which Azure SQL offering provides near-100 percent compatibility with on-premises Microsoft SQL Server instances, while providing automated updates, backups, and maintenance tasks?

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

Answer: A

Question: 243

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

Relational data uses to enforce relationships between different tables.

collections
rows
keys
partitions

Answer:

Answer Area

Relational data uses to enforce relationships between different tables.

collections
rows
keys
partitions

Question: 244

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

An extract load, and transform (ELT) process requires

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:

Answer Area

An extract load, and transform (ELT) process requires

- 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

Explanation:

A target data store powerful enough to transform data."

Explanation:

In an ELT (Extract, Load, Transform) process:

Data is first extracted from the source and then loaded into the target system (e.g., a data warehouse).

The transformation occurs after loading, leveraging the computational power of the target data store (such as a modern data warehouse or database that supports transformation operations directly).

Why the other options are incorrect:

"A data pipeline that includes a transformation engine":

This describes an ETL (Extract, Transform, Load) process, where transformation happens before loading the data into the target system.

"A separate transformation engine":

ELT relies on the target system to perform transformations, so a separate transformation engine is not required.

"Data that is fully processed before being loaded to the target data store":

This is again characteristic of ETL, where data is transformed before being loaded into the target system.

Important Tip:

ELT is best suited for modern cloud-based or high-performance data warehouses like Snowflake, BigQuery, or Redshift, where the system can handle large-scale transformations efficiently.

Question: 245

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

In an analytical model designed as a star schema, the entities by which you aggregate or slice measures are called

▼

cubes.

dimensions.

facts.

schemas.

Answer:

Answer Area

In an analytical model designed as a star schema, the entities by which you aggregate or slice measures are called

▼

cubes.

dimensions.

facts.

schemas.

Question: 246

CertyIQ

DRAG DROP

Match the data storage options to the appropriate types of data.

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

NOTE: Each correct match is worth one point.

Data storage options

Answer Area

Azure Blob Storage

Azure Cosmos DB for NoSQL

Azure Data Factory

Azure SQL Database

Audio files

JSON documents

Tabular datasets

Answer:

Data storage options

Answer Area

Azure Blob Storage

Azure Cosmos DB for NoSQL

Azure Data Factory

Azure SQL Database

Azure Blob Storage

Azure Cosmos DB for NoSQL

Azure SQL Database

Audio files

JSON documents

Tabular datasets

Explanation:

1. Azure Blob Storage → Audio files

- Azure Blob Storage is designed to store large unstructured data like audio files, video files, images, and logs.

2. Azure Cosmos DB for NoSQL → JSON documents

- Azure Cosmos DB is a globally distributed, NoSQL database that efficiently handles semi-structured data, including JSON documents.

3. Azure SQL Database → Tabular datasets

- Azure SQL Database is a relational database service that stores structured data in tabular format, making it ideal for relational datasets and SQL queries.

Each storage option in Azure is optimized for a specific type of data and use case:

Blob Storage handles unstructured data.

Cosmos DB is built for scalable NoSQL and JSON-based applications.

SQL Database is best for relational, structured, tabular datasets.

Important Tip:

Understand the purpose of Azure storage options to choose the correct service for your data storage and processing needs.

Question: 247

CertyIQ

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.

Answer Area

Statements	Yes	No
The Cool access tier is the default tier of an Azure Storage account.	<input type="radio"/>	<input type="radio"/>
Append blobs are optimized for random read and write operations.	<input type="radio"/>	<input type="radio"/>
Retrieving data from the Archive access tier takes hours to complete.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements	Yes	No
The Cool access tier is the default tier of an Azure Storage account.	<input type="radio"/>	<input checked="" type="radio"/>
Append blobs are optimized for random read and write operations.	<input type="radio"/>	<input checked="" type="radio"/>
Retrieving data from the Archive access tier takes hours to complete.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

The default access tier for a new general-purpose v2 storage account is set to the hot tier by default.

An append blob is composed of blocks and is optimized for append operations. When you modify an append blob, blocks are added to the end of the blob only, via the Append Block operation. Updating or deleting of existing blocks is not supported.

Rehydrating a blob from the Archive tier can take several hours to complete.

Question: 248

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

The Azure Cosmos DB

▼

for NoSQL
for Apache Gremlin
for MongoDB
for Table

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

Answer:

Answer Area

The Azure Cosmos DB

▼

for NoSQL
for Apache Gremlin
for MongoDB
for Table

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

Explanation:

The Azure Cosmos DB API for NoSQL supports querying documents using SQL. This article provides a sample document and two sample SQL queries and results / found in Microsoft tutorial

Question: 249

CertyIQ

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.

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:

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"/>

Question: 250

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

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:

Answer Area

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

Question: 251

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

After you create Microsoft Power BI data models and reports, you can publish them to

- Power BI Desktop.
- Power Query.
- the Power BI phone app.
- the Power BI service.

Answer:

Answer Area

After you create Microsoft Power BI data models and reports, you can publish them to

- Power BI Desktop.
- Power Query.
- the Power BI phone app.
- the Power BI service.

Question: 252

CertyIQ

Which type of Azure resource supports the serverless configuration of an Azure SQL database?

- A. SQL Server on Azure Virtual Machines
- B. an Azure SQL Database elastic pool
- C. a single database in Azure SQL Database
- D. Azure SQL Managed Instance

Answer: C

Explanation:

a single database in Azure SQL Database.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview?view=azuresql&tabs=general-purpose>

Question: 253

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

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

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

Answer:

Answer Area

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

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

Question: 254

CertyIQ

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.

Language

Answer Area

JavaScript

KQL

Python

SQL

Azure Data Explorer

Azure Stream Analytics

Azure Synapse Analytics

Answer:

Language

Answer Area

JavaScript

KQL

Azure Data Explorer

KQL

SQL

Azure Stream Analytics

Python

SQL

Azure Synapse Analytics

SQL

Explanation:

Azure Data Explorer: KQL (Kusto Query Language), which is a proprietary query language specifically designed for Azure Data Explorer. It is used to analyze and query large volumes of structured and unstructured data in real-time.

Azure Stream Analytics: SQL, which is a standard query language for relational databases. It is used to analyze and process high-velocity, real-time streaming data from various sources, such as IoT devices, social media platforms, and other real-time data streams.

Azure Synapse Analytics: SQL.

Question: 255

CertyIQ

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

- A. Table
- B. Cassandra
- C. NoSQL
- D. Gremlin

Answer: D

Explanation:

Correct answer is D:Gremlin.

Question: 256

CertyIQ


HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

A Microsoft Power BI

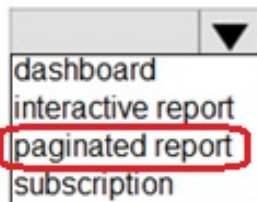


enables users to create highly formatted, fixed-layout documents optimized for printing and archiving.

Answer:

Answer Area

A Microsoft Power BI



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

Explanation:

The Power BI feature that allows users to create highly-formatted, fixed layout documents optimized for printing and archiving is called "Power BI Paginated Reports" (previously known as SQL Server Reporting Services or SSRS).

Power BI Paginated Reports are designed for creating pixel-perfect, paginated documents that can be printed or exported to PDF, Word, or other file formats. Paginated Reports offer a fixed layout and allow for precise control over the formatting of data, making them ideal for creating documents like invoices, financial statements, and other documents that require consistent formatting.

With Paginated Reports, users can connect to a variety of data sources, including SQL Server, Azure SQL Database, Oracle, and many others, and can create rich visualizations and reports that can be easily shared and printed. Paginated Reports can be created using either Power BI Report Builder or through the Power BI service.

Question: 257

CertyIQ

What is a difference between structured data and semi-structured data?

- A. Structured data has a fixed schema and semi-structured data has a flexible schema.
- B. Only Structured data supports entities.
- C. Structured data has a flexible schema and semi-structured data has a fixed schema.
- D. Only Structured data supports attributes.

Answer: A

Explanation:

Whereas structured data requires a fixed schema defined in advance, semi-structured data does not require a

prior schema definition. For this reason, it can constantly evolve — new attributes can be added at any time. So therefore A is correct answer.

Question: 258

CertyIQ

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.

Answer Area

Statement

Yes

No

Database administrators apply data cleansing routines and turn data into useful information.

☐☐

Data engineers manage databases, store backup copies of data, and restore data in the event of a failure.

☐☐

Data analysts create data visuals and enable companies to make data-driven decisions.

☐☐

Answer:

Answer Area

Statement

Yes

No

Database administrators apply data cleansing routines and turn data into useful information.

☐☒

Data engineers manage databases, store backup copies of data, and restore data in the event of a failure.

☐☒

Data analysts create data visuals and enable companies to make data-driven decisions.

☒☐

Explanation:

no

no

yes

Question: 259

CertyIQ

A bank needs to ensure that after an account transfer transaction completes, the revised account balances persists even if the database system hosting the transaction becomes temporarily unavailable.

Of which ACID semantic is this an example?

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

Answer: A

Explanation:

durability is a correct answer.

Question: 260

CertyIQ

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

- A.INSERT
- B.ALTER
- C.DROP
- D.CREATE

Answer: A

Explanation:

Correct answer is A:INSERT

Question: 261

CertyIQ

Which language is used to define queries in Azure Synapse Data Explorer?

- A.Bash
- B.PowerShell
- C.KQL
- D.SQL

Answer: C

Explanation:

C.KQL

Question: 262

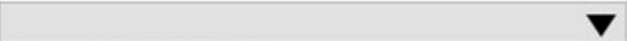
CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

The  service is optimized for IoT solutions.

- Azure SQL Database
- Azure SQL Edge
- Azure SQL Managed Instance
- SQL Server on Azure Virtual Machines

Answer:

Answer Area

The  service is optimized for IoT solutions.

- Azure SQL Database
- Azure SQL Edge**
- Azure SQL Managed Instance
- SQL Server on Azure Virtual Machines

Question: 263

CertyIQ

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.

Answer Area

Statement	Yes	No
In a relational database, the schema is defined before data is written to the table.	<input type="radio"/>	<input type="radio"/>
Key/value data stores are highly optimized for complex queries that contain multiple joins.	<input type="radio"/>	<input type="radio"/>
A relational database is commonly used to store and query structured data.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statement	Yes	No
In a relational database, the schema is defined before data is written to the table.	<input checked="" type="radio"/>	<input type="radio"/>
Key/value data stores are highly optimized for complex queries that contain multiple joins.	<input type="radio"/>	<input checked="" type="radio"/>
A relational database is commonly used to store and query structured data.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

yes

no

yes

Question: 264

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

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:

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

Question: 265

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

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

CSV

JSON

ORC

TSV

Answer:

Answer Area

The

CSV
JSON
ORC
TSV

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

Question: 266

CertyIQ

DRAG DROP

Match the Azure Cosmos DB APIs to the appropriate data structures.

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

NOTE: Each correct match is worth one point.

APIs

Answer Area

Azure Cosmos DB for
Apache Cassandra API

Azure Cosmos DB for
Apache Gremlin API

Azure Cosmos DB for
NoSQL API

Azure Cosmos DB for
Table API

API

Graph data

API

JSON documents

API

Key/value data

Answer:

Answer Area

Azure Cosmos DB for
Apache Gremlin API

Graph data

Azure Cosmos DB for
NoSQL API

JSON documents

Azure Cosmos DB for
Table API

Key/value data

Question: 267

CertyIQ

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

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

Answer: D

Explanation:

Azure Cosmos DB for Apache Gremlin

Question: 268

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

▼

Data analysts

Data engineers

Data scientists

Database administrators

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

Answer:

Answer Area

Data analysts

Data engineers

Data scientists

Database administrators

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

Explanation:

Correct answer is Data base administrators.

Question: 269

CertyIQ

DRAG DROP

Match the cloud service model to the appropriate Azure relational database service.

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

NOTE: Each correct match is worth one point.

Cloud service model

Infrastructure as a service

Platform as a service

Software as a service

Answer Area

Cloud service model

Azure SQL Database

Cloud service model

Azure SQL Managed Instance

Cloud service model

SQL Server on Azure Virtual Machines

Answer:

Answer Area

Platform as a service

Azure SQL Database

Platform as a service

Azure SQL Managed Instance

Infrastructure as a service

SQL Server on Azure Virtual Machine:

Explanation:

Platform as a service.

Platform as a service.

Question: 270

CertyIQ

A bank needs to ensure that a transaction involving debiting funds from a source account and crediting the same funds to a destination account must complete both actions. If either action falls to complete, the other action must fail.

Of which ACID semantic is this an example?

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

Answer: C

Explanation:

Atomicity guarantees that each transaction is treated as a single "unit", which either succeeds completely or fails completely: if any of the statements constituting a transaction fails to complete, the entire transaction fails and the database is left unchanged. Consistency ensures that a transaction can only bring the database from one consistent state to another, preserving database invariants: any data written to the database must be valid according to all defined rules. Isolation ensures that concurrent execution of transactions leaves the database in the same state that would have been obtained if the transactions were executed sequentially. Durability guarantees that once a transaction has been committed, it will remain committed even in the case of a system failure (e.g., power outage or crash).

Question: 271

CertyIQ

Blob rehydration occurs in Azure Blob storage when a blob moves between which access tiers?

- A.Hot to Cool
- B.Hot to Archive
- C.Cool to Archive
- D.Archive to Cool

Answer: D

Explanation:

Correct answer is D:Archive to Cool.

Reference:

<https://learn.microsoft.com/en-us/azure/storage/blobs/archive-rehydrate-overview>

Question: 272

CertyIQ

Which activity is most common for transactional workloads?

- A.recording small units of work events in real time
- B.aggregating massive amounts of data
- C.producing complex reports
- D.performing self-service data exploration

Answer: A

Explanation:

A. recording small units of work events in real time.

Question: 273

CertyIQ

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

Answer: A

Explanation:

Correct: A. online transaction processing (OLTP).

Question: 274

CertyIQ

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.

Job roles

Data analyst

Data architect

Data engineer

Database administrator

Answer Area

Job role

Create a pipeline to transfer data from a database to a data warehouse.

Job role

Create a transactional database.

Job role

Create a chart that shows sales transactions by week.

Answer:

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

Explanation:

Data engineer

Database administrator

Data analyst.

Question: 275

CertyIQ

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.

Job role	Answer Area
Data analyst	Job role Database backup and recovery
Data engineer	Job role Enabling advanced analytic capabilities through reports
Database administrator	Job role Monitoring data pipelines

Answer:

Answer Area

Database administrator

Database backup and recovery

Data analyst

Enabling advanced analytic capabilities through reports

Data engineer

Monitoring data pipelines

Explanation:

Database administrator

Data analyst

Data engineer.

Question: 276

CertyIQ

DRAG DROP

Match the Azure Data Lake Storage Gen 2 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.

Terms

Azure Storage account

File share

Container

Answer Area

Azure Resource Group

Folders

Files

Answer:

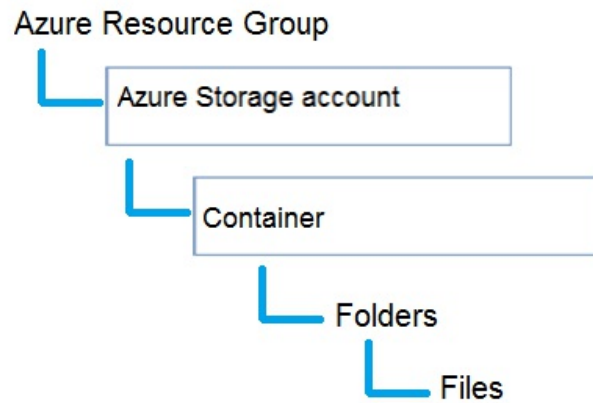
Terms

Azure Storage account

File share

Container

Answer Area



Explanation:

Box 1: Azure Storage account -

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

Box 2: Container -

Reference:

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

Question: 277

CertyIQ

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.

Job role

Data analyst

Data engineer

Database administrator

Answer Area

Assign permissions to users.

Derive meaning from data by creating visualizations and charts.

Implement which pipelines to transfer and transform data between systems.

Answer:

Answer Area

Database administrator

Assign permissions to users.

Data analyst

Derive meaning from data by creating visualizations and charts.

Data engineer

Implement which pipelines to transfer and transform data between systems.

Explanation:

Database administrator

Data analyst

Data engineer.

Question: 278**CertyIQ**

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

	▼
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:**Answer Area**

	▼
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.

Explanation:

Correct answer is :A data analyst.

Question: 279**CertyIQ**

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)

Answer: C**Explanation:**

Data Manipulation Language (DML).

Question: 280**CertyIQ**

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

▼

CREATE
DROP
RENAME
UPDATE

is an example of a Data Manipulation Language (DML) statement.

Answer:

Answer Area

▼

CREATE
DROP
RENAME
UPDATE

is an example of a Data Manipulation Language (DML) statement.

Question: 281

CertyIQ

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.

Answer Area

Statements	Yes	No
The pgAdmin tool can be used to manage Azure Database for PostgreSQL.	<input type="radio"/>	<input type="radio"/>
Azure Database for MariaDB offers point-in-time restore for up to 365 days.	<input type="radio"/>	<input type="radio"/>
Azure Database for MySQL always requires SSL connections.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements

	Yes	No
The pgAdmin tool can be used to manage Azure Database for PostgreSQL.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Database for MariaDB offers point-in-time restore for up to 365 days.	<input type="radio"/>	<input checked="" type="radio"/>
Azure Database for MySQL always requires SSL connections.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:

Yes

No

No

Question: 282

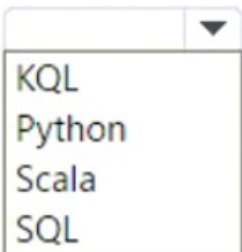
CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.


Answer Area

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

KQL
Python
Scala
SQL

Answer:

Answer Area

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

KQL
Python
Scala
SQL

Explanation:

Correct answer is KQL.

Question: 283

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

An online analytical processing (OLAP) model stores data that is optimized for analytical workloads.

aggregated
normalized
semi-structured
unstructured

Answer:

Answer Area

An online analytical processing (OLAP) model stores data that is optimized for analytical workloads.

aggregated
normalized
semi-structured
unstructured

Explanation:

Correct answer is A:Aggregated.

Question: 284

CertyIQ

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

Answer: C

Explanation:

A Fact table typically stores quantitative data, often referred to as measures or metrics, and contains foreign keys that reference related Dimension tables. Fact tables are central to a star schema or snowflake schema design, which are common schema designs used in data warehousing and OLAP environments for supporting analytical querying and reporting.

Question: 285

CertyIQ

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.

Answer: A

Explanation:

Data is read from a single table and written to a single table.

Question: 286

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

In Microsoft Power BI, you can create

a dashboard
a fact table
a unique key
an attribute hierarchy

to quickly drill-up or drill-down to find aggregated values at different levels of a dimension.

Answer:

Answer Area

In Microsoft Power BI, you can create

a dashboard
a fact table
a unique key
an attribute hierarchy

to quickly drill-up or drill-down to find aggregated values at different levels of a dimension.

Explanation:

an attribute hierarchy. An attribute hierarchy allows you to organize dimension attributes into a hierarchical structure, typically with multiple levels of granularity. This enables users to easily navigate through different levels of detail within a dimension when analyzing data.

Question: 287

CertyIQ

HOTSPOT

-

Select the answer that correctly completed the sentence.

Answer Area

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

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

Answer:

Answer Area

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

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

Explanation:

Correct answer is D:Azure synapse analytics.

Question: 288

CertyIQ

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

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

Answer:

Answer Area

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

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

Explanation:

azure synapse analytics is correct answer.

Question: 289

CertyIQ

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

▼

a cube.
an aggregated dimension.
an aggregated measure.

Answer:

Answer Area

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

▼

a cube.
an aggregated dimension.
an aggregated measure.

Explanation:

The answer a Cube.

Reference:

<https://learn.microsoft.com/en-us/analysis-services/multidimensional-models/cubes-in-multidimensional-models?view=asallproducts-allversions>

Question: 290

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

▼

Azure Databricks
Azure HDInsight
Azure Stream Analytics
Azure Synapse Analytics

combines a scalable, high-performance Microsoft SQL Server-based relational data warehouse

with the flexibility of a data lake and open-source Apache Spark.

Answer:

Answer Area

▼

- Azure Databricks
- Azure HDInsight
- Azure Stream Analytics
- Azure Synapse Analytics**

combines a scalable, high-performance Microsoft SQL Server-based relational data warehouse

with the flexibility of a data lake and open-source Apache Spark.

Explanation:

Correct answer is Azure Synapse Analytics.

Question: 291

CertyIQ

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

Azure Blob Storage page blobs are optimized for

▼

- append operations.
- random read and write operations.
- storing key/value data.
- storing large binary objects that change infrequently.

Answer:

Answer Area

Azure Blob Storage page blobs are optimized for

▼

- append operations.
- random read and write operations.**
- storing key/value data.
- storing large binary objects that change infrequently.

Explanation:

random read and write operations.

Question: 292

CertyIQ

GRANT, REVOKE, and DENY are examples of which type of SQL statement?

- A.Data Definition language (DDL)
- B.Data Control Language (DCL)

Answer: B

Explanation:

Data Control Language (DCL).

Question: 293

CertyIQ

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 all.

NOTE: Each correct match is worth one point.

Terms

Answer Area

Atomicity

Concurrent transactions cannot interfere with one another and must result in a consistent database state.

Consistency

Each transaction is treated as a single unit that succeeds completely or fails completely.

Isolation

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.

Answer:

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.

Question: 294

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.

Terms

Attributes

Entities

Foreign keys

Primary keys

Schemas

Answer Area

Get separated into individual columns

Get separated into individual tables

Link two related entities

Answer:

Answer Area

Attributes

Get separated into individual columns

Entities

Get separated into individual tables

Foreign keys

Link two related entities

Explanation:

Attribute - Get separated into individual columns.

In a database, an attribute typically refers to a column in a table, which is where individual data elements of an entity are stored.

Entities - Get separated into individual tables.

Entities in a database are typically represented as tables. Each table represents a specific type of object or concept within the system.

Foreign Keys - Link two related entities.

A foreign key is a field (or collection of fields) in one table that uniquely identifies a row of another table. It is used to establish and enforce a link between the data in two tables.

Question: 295**CertyIQ**

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

To create a dashboard, you can use

- only the Microsoft Power BI Desktop app
- only the Microsoft Power BI mobile app
- only the Microsoft Power BI service
- either the Microsoft Power BI Desktop app or the Power BI service

Answer:**Answer Area**

To create a dashboard, you can use

- only the Microsoft Power BI Desktop app
- only the Microsoft Power BI mobile app
- only the Microsoft Power BI service**
- either the Microsoft Power BI Desktop app or the Power BI service

Explanation:

Only the Microsoft Power BI service.

Dashboards can only be created in Power BI Service (Online), not in Power BI Desktop or the Mobile App.

Power BI Desktop is used to design reports, and these reports can then be used in Power BI Service to build dashboards.

Question: 296**CertyIQ**

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.

Answer Area**Statements****Yes****No**

Batch processing is optimized for handling large datasets.

☐☐

Stream processing is optimized for performing complex analysis of datasets.

☐☐

The latency of batch processing is typically measured in seconds or milliseconds.

☐☐

Answer:

Answer Area

Statements

Yes

No

Batch processing is optimized for handling large datasets.

☒☐

Stream processing is optimized for performing complex analysis of datasets.

☐☒

The latency of batch processing is typically measured in seconds or milliseconds.

☐☒

Explanation:

Yes

No

No

Question: 297

CertyIQ

Structured data where each row represents a single data entity uses which type of schema?

- A.XML
- B.tabular
- C.blob
- D.JSON

Answer: B

Explanation:

Correct answer is B:tabular.

Question: 298

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Answer Area

The workflow for creating an analytics solution in Microsoft Power BI starts with

Power BI Desktop.
the Power BI phone app.
the Power BI service.

Answer:

Answer Area

The workflow for creating an analytics solution in Microsoft Power BI starts with

Power BI Desktop.
the Power BI phone app.
the Power BI service.

Question: 299

CertyIQ

You need to implement an Azure platform as a service (PaaS) service that will host a relational database. The solution must support built-in autoscaling.

Which service should you implement?

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

Answer: A

Explanation:

Correct answer is A: Azure SQL Database.

Question: 300

CertyIQ

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.

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:

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"/>

Explanation:

No

No

Yes

Question: 301

CertyIQ

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.

Answer Area

Statements	Yes	No
Each Azure SQL managed instance supports multiple databases that can be accessed by using cross-database queries.	<input type="radio"/>	<input type="radio"/>
Elastic pools are used to share resources across multiple instances of SQL Server on Azure Virtual Machines.	<input type="radio"/>	<input type="radio"/>
Azure SQL Database is fully compatible with both on-premises physical instances and virtualized instances of Microsoft SQL Server.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements	Yes	No
Each Azure SQL managed instance supports multiple databases that can be accessed by using cross-database queries.	<input checked="" type="radio"/>	<input type="radio"/>
Elastic pools are used to share resources across multiple instances of SQL Server on Azure Virtual Machines.	<input checked="" type="radio"/>	<input type="radio"/>
Azure SQL Database is fully compatible with both on-premises physical instances and virtualized instances of Microsoft SQL Server.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:

yes

yes

No

Question: 302

CertyIQ

You have the following data.

```
{
  "firstName": "John",
  "lastName": "Doe",
  "address":
  {
    "streetAddress": "1 Main St.",
    "city": "New York",
    "state": "NY",
    "postalCode": "10099"
  }
},
{
  "firstName": "Jane",
  "lastName": "Doe",
  "address":
  {
    "streetAddress": "1 Main St.",
    "city": "New York",
    "state": "NY",
    "postalCode": "10099"
  },
  "contact":
  [
    {
      "type": "email",
      "address": "jane.doe@contoso.com"
    }
  ]
}
```

Which type of data is this?

- A.unstructured
- B.semi-structured
- C.structured

Answer: B

Explanation:

Correct answer is B:semi-structured.

Thank you

Thank you for being so interested in the premium exam material.
I'm glad to hear that you found it informative and helpful.

If you have any feedback or thoughts on the bumps, I would love to hear them.
Your insights can help me improve our writing and better understand our readers.

Best of Luck

You have worked hard to get to this point, and you are well-prepared for the exam
Keep your head up, stay positive, and go show that exam what you're made of!

Feedback

More Papers



Future is Secured
100% Pass Guarantee



24/7 Customer Support
Mail us - certyiqofficial@gmail.com



Free Updates
Lifetime Free Updates!

Total: **302 Questions**

Link: <https://certyiq.com/papers/microsoft/dp-900>