



"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks.com

Warning: Keep connected with our support team
for latest updates

Question: 1

HOTSPOT

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

Answer Area

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

ICMngM lo rsuioimhipi between uMw" geography illy drstnbUed
writer hfafUMjicrjwnfn

Answer:

Explanation:

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

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: 2

HOTSPOT

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

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 processed before being loaded to the target data store

Answer

Explanation:

Answer Area

An extract, load, and transform (ELT) process requires data that is fully processed before being loaded to the target data store.

Question: 3

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

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

Answer: B

Explanation:

Question: 4

HOTSPOT

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

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

Answer:

Explanation:

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

cognitive descriptive predictive
prescriptive

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: 5

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. Each location may be used 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.

Locations	Answer Area
An in-memory data integration tool	Extract: Location
The CRM system	Load: Location
The data warehouse	Transform: Location

Answer:

Extract:

The CRM system

Load:

The data warehouse

Transform:

An in-memory data integration tool

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: An in-memory data integration tool

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




Reference:

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

Question: 6

Your company recently reported sales from the third quarter.

You have the chart shown in the following exhibit.



Which type of analysis is shown for the fourth quarter?

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

Answer: A

Explanation:

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

DRAG DROP

Match the types of workloads the appropriate scenario.

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.

Workload Type*

Batch

Micro-batch

Streaming

Answer Area

Workload Type Data for a product catalog will be loaded every 12

hours to a data warehouse

Workload Type Data for online purchases will be loaded to a

data warehouse as the purchases occur.

Workload Type Updates to inventory data will be loaded to a data

warehouse every 1,000 transactions

Answer:

Explanation:

Workload Type*

Batch

Micro-batch

Streaming

Answer Area

Batch

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

Streaming

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

Micro-batch

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

Question: 8

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 can output data to a data store

Batch processing can output data to a relational database

Batch processing can output data to a NoSQL database

Answer:

Explanation:

Answer Art*

Statements

Yes

No

Batch processing can output data to a file store

Batch processing can output data to a relational database

Batch processing can output data to a NoSQL database

Question: 9

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.

Analytics Types

Cognitive

Answer Arcs

Diagnostic

Analytical Type How do I allocate my budget to buy different inventory items?

Analytical Type Which people are mentioned in a company's business documents?

Answer:

Explanation:

Diagnostic

Why did sales increase last month?

Predictive

How do I allocate my budget to buy different inventory items?

Cognitive

Which people are mentioned in a company's business documents?

Reference:

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

Question: 10

HOTSPOT

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

Answer

Are*

Relational data is stored in

rows and columns

rows and columns

rows and columns

rows and columns

rows and columns

Answer:

Explanation:

Answer Are*

ReUtlomi dan d stored in a hwrarcti.ii fotdw rtmotur*

Question: 11

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

NOTE: Each correct selection is worth one point.

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

Answer: A, B

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: 12

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.



Answer Area

Workload Type Analysing web Traffic from the past year

Workload Type Classifying images that were uploaded last month
Workload Type tracking how many people are visiting a website


Answer:

service may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.


Location1
Azure Analysis Azure Colima
Services | DB
AzmeDat# 1| Azure Synapse
factory | Analytics
Anw Table storage

Answer Area



Answer:

Explanation:



Question: 15

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 (OLTP)
- D. Online Transaction Processing (OLTP)

Answer: D

Explanation:

Reference:

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

Question: 16

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

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

Answer: AB

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: 17

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

NOTE: Each correct selection is worth one point.

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

Answer: A, D

Explanation:

Reference:

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

Question: 18

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

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

Answer: D

Explanation:

Reference:

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

<https://docs.microsoft.com/en-us/power-bi/visuals/power-bi-visualization-kpi>

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

Question: 19

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

- A. Power BI report Builder
- B. Power BI Desktop
- C. Charticulator
- D. the Power BI service

Answer: A

Explanation:

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

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

Question: 20

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

Yes

No

A pipeline is a representation of a On* structure within Azure Data Factory

Azure Data Factory pipelines can execute other pipelines

A processing step within an Azure Data Factory pipeline is an activity.

Answer:

Explanation:

Answer Area

Statements

Yes

No

A pipeline is a representation of a data structure within Azure Data

Factory

0

Azure Data Factory pipelines can execute other pipelines.

A processing step within an Azure Data Factory pipeline is an activity.

0

4

Question: 21

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.

Component*

0/12

1 Linked service 11 Mapping data flow Pipeline

Answer Area

Component A represents data structures within data stores

Component B The information used to connect to external resources

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

Answer:

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

Box 1: Dataset

Datasets must be created from paths in Azure data stores 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: 22

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

Answer: B

Explanation:

Reference:

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

Question: 23

DRAG DROP

Match the Azure services to appropriate requirements.

To answer, drag the appropriate services 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.

Azure Services

Acute Oat* Factory

I Azure Oats lake Storage

Azure MI Database

Azure Synapse Analytics

Answer Area

Azure Service

Output daw to Parquet fomwt

Azure Service

Stare data that is in Parquet formal

Azure Service

Pewit a uMir riptNntMon of su that li stated
In Parquet formal

Answer:

Explanation:

Azure Services



Answer Area

Azure Data Factory	Output data to Parquet format.
Azure Data Lake Storage	Store data that is in Parquet format.
Azure Synapse Analytics	Persist a tabular representation of data that is stored in Parquet format.

Question: 24

HOTSPOT

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

Answer Area

of Azure Synapse Analytics
 of Azure Synapse Analytics
 of Azure Synapse Analytics
 of Azure Synapse Analytics

Answer

Explanation:

Answer Area

The multi-tenant parallel processing (MPP) engine
 of Azure Synapse Analytics redirects client connections across compute nodes.

Question: 25

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

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

Answer: B

Explanation:

Reference:

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

Question: 26

You have a transactional application that stem 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.

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

Question: 27

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
Azure Synapse Analytics scales storage and compute independently	☐	☐
Azure Synapse Analytics can be paused to reduce compute costs	☐	☐
An Azure Synapse Analytics data warehouse has a fixed storage capacity	☐	☐

Answer:

Explanation:

Answer Area

Statements	Yes	No
Azure Synapse Analytics scales storage and compute independently	☒	☐
Azure Synapse Analytics can be paused to reduce compute costs	☒	☐
An Azure Synapse Analytics data Warehouse has a fixed storage capacity	☐	☒

Question: 28
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.

Statements	Yes	No
Ann Ottoback on consume dtc* from Azure SQL Detsbest	☐	☒
Azure OjUbncks cm consume dtc* from Azure Event Hubs.	☐	☐
Azure OIUbricki cm consume dtc* from Azure Cosmos DB	☐	☐

Answer:

Explanation:


Question: 29

HOTSPOT

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

Answer Area

A key/value data store is optimized for



	▼
enforcing constraints	
simple lookups	
table joins	
transactions	

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: 30

You have an application that runs on Windows and requires access to a mapped drive.

Which Azure service should you use?

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

Answer: A

Explanation:

Question: 31

HOTSPOT

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

Answer Area

You can query a graph database in Azure Cosmos DB as a JSON document by using a SQL-like language.



Answer:

Explanation:

Answer Area

You can query a graph database in Azure Cosmos DB as a JSON document by using a SQL-like language.

Question: 32

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

Which type of data store should you use?

- A. Key/value
- B. Graph
- C. Document
- D. columnar

Answer: B

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: 33

You need to recommend a non-relational data store that is optimized for storing and retrieving files,

videos, audio stream, 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. document
- B. key/value
- C. object
- D. columnar

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: 34

You need to create an Azure Storage account.
Data in the account must replica outside the Azure region automatically.

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

NOTE: Each correct selection is worth one point.

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

Answer: A, C

Explanation:

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

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

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: 35

HOTSPOT

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

Answer Area

Statements	Yes	No
The Azure Cosmos DB API is configured separately for each database in an Azure Cosmos DB account.	☐	☐
Partition keys are used in Azure Cosmos DB to optimize queries.	☐	☐
Items contained in the same Azure Cosmos DB logical partition can have different partition keys.	☐	☐

Answer:

Explanation:

Answer Area

Statements	Yes	No
The Azure Cosmos DB API is configured separately for each database in an Azure Cosmos DB account.	☐	☐
Partition keys are used in Azure Cosmos DB to optimize queries.	☐	☐
Items contained in the same Azure Cosmos DB logical partition can have different partition keys.	☐	☐

Question: 36

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 automate the creation of an interdependent and develop can deploy from the Azure portal
- C. To provision Azure subscriptions and apply policies for multi-tenant deployments
- D. To set complex permissions automatically on existing objects in an Azure subscription

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: 37

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

- A. multi-master replication

- B. Availability Zones
- C. automatic failover
- D. the strong consistency level

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.

Reference:

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

Question: 38

At which level in Azure Cosmos DB can you configure multiple write regions and read regions?

- A. database
- B. partition
- C. collection
- D. account

Answer: C

Explanation:

Question: 39

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

Answer: B

Explanation:

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

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: 40

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. container
- B. item
- C. database
- D. partition

Answer: AC

Explanation:

Reference:

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

Question: 41

DRAG DROP

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

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

NOTE: Each correct match is worth one point.

Azure Storage account


1 Hectare

Storage

Level

Answer:

Explanation:



Box 1: Azure Storage account

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

Box 2: File share

Reference:

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

Question: 42

HOTSPOT

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

Answer Area

To configure an Azure Storage account to support both security at the folder level and atomic directory manipulation, enable the `fileShares` namespace set account kind to `FileShare` set Performance to Premium

Answer:

Explanation:

Answer Area

To configure an Azure Storage account to support both security at the folder level and atomic directory manipulation, enable the `fileShares` namespace set Account kind to Blob Storage

Question: 43

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. columnar
- C. graph
- D. time series

Answer: D

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: 44

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

- A. supports partitioning
- B. provides resiliency if any Azure region fails
- 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: 45

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: 46

DRAG DROP

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

To answer, drag the appropriate data types 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.

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:

Explanation:

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

Box 1: Image files

Azure Blob storage is suitable for image files.

Box 2: Key/value pairs

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

Box 3: Relationship between employees

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

Reference:

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

Question: 47

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.

NOTE: Each correct match is worth one point.

Tool	Answer Area
Alme Data Studio	Tool A graphical tool for managing Azure SQL databases and viewing SQL execution plans
Microsoft ML Sew Data TmfcgSm	Tool A Microsoft SQL Server extension that supports connections to SQL Server and provides a rich editing experience for T-SQL
Microsoft SQL Sew Management Studio SSMS1	Tool A lightweight editor that can run on-demand SQL queries and view and save results as text, JSON, or Microsoft Excel files
Microsoft Visual Studio Code	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:

Explanation:

Microsoft SQL Server Management Studio	A graphical tool for managing SQL Server or Azure SQL databases that supports access, configuration, management, and administration tasks
--	---

Microsoft Visual Studio Code

A lightweight source code editor with an mssql extension that supports connections to SQL Server and a rich editing experience for T-SQL

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 relational databases, SQL Server Analysis Services (SSAS) data models, SQL Server Integration Services (SSIS) packages, and SQL Server Reporting Services (SSRS) reports.

Question: 48

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

- A. INSERT
- B. DISABLE
- C. ALTER
- D. DROP

Answer: A

Explanation:

Data Manipulation Language (DML) statements:

DELETE
INSERT
UPDATE

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-reference-tsql-statements>

Question: 49

You have a SQL query that combines customer data and order data.

a. The query includes calculated columns.

You need to persist the SQL query so that other users can use the query. What should you create?

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

Answer: D

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: 50

HOTSPOT


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

Answer Are#



Answer:

Explanation:



Reference:

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

Question: 51

HOTSPOT

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

NOTE Each correct selection is worth one point.

Answer Are* Statements

Yes

No

Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than infrastructure as a service (IaaS) database offerings.
Platform as a service (PaaS) database offerings in Azure provide administrators with the ability to control and update the operating system version.
All platform as a service (PaaS) database offerings in Azure can be paused to reduce costs. 0

Answer:

Explanation:

Answer Area

Statements

Yes

No

Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than infrastructure as a service (IaaS) database offerings
Platform as a service (PaaS) database offerings in Azure provide administrators with the ability to control and update the operating system version
Azure platform as a service (PaaS) database offerings in Azure can be paused to reduce costs

Question: 52

Which Azure SQL offering supports automatic database scaling and automatic pausing of the database during inactive periods?

- A. Azure SQL Database Hyperscale
- B. Azure SQL managed instance
- C. Azure SQL Database serverless
- D. Azure SQL Database elastic pool

Answer: B

Explanation:

Question: 53

HOTSPOT

You have the following SQL query.

```
INSERT INTO dbo.Product (ProductID, ProductName, ProductDescription)
VALUES (1, 'Lamp', U.S. 'Bulb', 'Light')
```

NOTE: Each correct selection is worth one point.

Answer Area

Dbo.Products:	☐ A column ☐ A database ☐ A table ☐ An index
ProductName:	☐ A column ☐ A database ☐ A table ☐ An index

Answer:

Explanation:

Dbo.Products:

A column
A database
A table
An index

ProductName:

A column
A database
A table
An index

Question: 54

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements

Relational database tables contain columns and rows

Indexes in a relational database describe the data types in a table

A database view is a virtual table whose content is defined by a query

Yes

No

Q

Q

0

0

0

0

Answer:

Explanation:

Statements

Yes

No

Relational database tables contain columns and rows

Q

Q

Indexes in a relational database describe the data types in

a table

0

0

A database view is a virtual table whose content is defined by a query

0

Box 1: Yes

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

Box 2: No

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

Box 3: Yes

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

Reference:

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

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

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

Question: 55

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

Answer: B, C

Explanation:

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.

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: 56

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

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

Answer: C

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: 57

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 Synapse Analytics
- C. Azure HDInsight
- D. Azure SQL Database

Answer: D

Explanation:

Azure SQL Database is relational database and a managed service.

Reference:

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

Question: 58

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.

Amwet Are#

Statements	Yes	No
If you have a platform as a service (PaaS) database in Azure, you are responsible for applying operating system updates.	☐	☐
If you have a platform as a service (PaaS) database in Azure, backups are performed automatically.	☐	☐
If you have a platform as a service (PaaS) database in Azure, you are responsible for upgrading the database engine.	☐	☐

Answer:

Explanation:

Statements	Yes	No
If you have a platform as a service (PaaS) database in Azure, you are responsible for applying operating system updates.	☐	☐
If you have a platform as a service (PaaS) database in Azure, backups are performed automatically.	☐	☐
If you have a platform as a service (PaaS) database in Azure, you are responsible for upgrading the database engine.	☐	☐

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: 59

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

You recently changed your external IP address.

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

Answer: A

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: 60

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

- A. the Azure Active director/ (Azure AO) account that created the database
- B. the Azure Active Directly Azure AD) administrator account
- C. the sa account
- D. the server admin login account of the logical server

Answer: D

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: 61

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
Azure Data Studio can be used to Query an Azure SQL database from a device that runs macOS.	☐	☐
Microsoft SQL Server Management Studio (SSMS) enables users to create and use SQL notebooks.	☐	☐
Azure Data Studio can be used to restore a database.	☐	☐

Answer:

Explanation:

Answer Area

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

Question: 62

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

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

Answer: A

Explanation:

Reference:

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

Question: 63

DRAG DROP

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

SalesDate	SalesAmount	Product ID
14-Apr-20	\$2000	1
5-Apr-20	\$40	2
5-Apr-20	\$2300	1
6-Apr-20	\$40	
1 Apr '20	\$200	4

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

SaleDate	AVG(SalesAmount)
4-Apr-20	\$2000
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. Values Answer Area

← OMUT: Mud

SW (L-700), W0, L^1

run select Value SriOrtr

Answer:

```
SELECT SalesDate, AVG(SalesAmount)
FROM Sales
GROUP BY SalesDate
ORDER BY SalesDate
```

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: 64

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
Azure SQL managed instance supports cross-database queries	☐	☐
Azure SQL managed instance supports user-created backups	☐	☐
Azure SQL managed instance databases can be restored to SQL Server on an Azure virtual machine	☐	☐

Answer:

Explanation:

Answer Area	Statements	Yes	No
	Azure SQL managed instance supports cross-database queries	☐	☐
	Azure SQL managed instance supports user-created backups	☐	☐
	Azure SQL managed instance databases can be restored to SQL Server on an Azure virtual machine	☐	☐

Question: 65

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 the computers. The developers require code assistance features such as IntelliSense.

What should the developers use?

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

Answer: A

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: 66

HOTSPOT

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

Answer Area

Transparent Data Encryption (TDE) encrypts _____ a column to protect data at rest and in transit queries and We* results in o^er to protect data m transit the database to prosed data at rest the wuel i? protect pm at ten

Answer:

Explanation:

Answer Area

Transparent Data Encryption (TDE) encrypts the database to protect data at rest.

Question: 67
HOTSPOT

To complete the sentence, select the appropriate option in the 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:

Explanation:

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.

Reference:

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

Question: 68
HOTSPOT

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

NOTE: Each correct selection is worth one point.

Statements

Yes

No

Normalization involves eliminating relationships between database tables.

☐☐

Normalizing a database reduces data redundancy.

☐☐

Normalization improves data integrity.

☐☐

Answer:

Explanation:

Statements

Yes

No

Normalization involves eliminating relationships between database tables.

☐☐

Normalizing a database reduces data redundancy.

☐☐

Normalization improves data integrity.

☐☐

Reference:

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

Question: 69

HOTSPOT

To complete the sentence, select the appropriate option in the 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:

Answer:

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.

Reference:

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

Question: 70

HOTSPOT

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

In batch processing,

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

Answer:

Explanation:

In batch processing,

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

Reference:

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

Question: 71

HOTSPOT

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

Transcribing audio files is an example of _____ analytics.

	▼
cognitive	
descriptive	
predictive	
prescriptive	

Answer:

Explanation:

Transcribing audio files is an example of _____ analytics.

	▼
cognitive	
descriptive	
predictive	
prescriptive	

Reference:

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

Question: 72

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" E (  
      "service" : "linkedin", "handle" :  
        "bensmith" }  
  E  
  "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.

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:

Explanation:

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	

Reference:


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

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

Question: 73

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

The data model is a [answer choice].

	▼
transactional model	
star schema	
snowflake schema	

Customer is a [answer choice] table.

	▼
fact	
dimension	
bridge	

Answer:

Explanation:

The data model is a [answer choice].

	▼
transactional model	
star schema	
snowflake schema	
▼	
fact	
dimension	
bridge	

Customer is a [answer choice] table.

Reference:

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

<https://azure.microsoft.com/en-us/blog/data-models-within-azure-analysis-services-and-power-bi/>

Question: 74

HOTSPOT

Answer:

To complete the sentence, select the appropriate option in the 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:

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.

Reference:

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

Question: 75

HOTSPOT

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

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

A clustered index

A FileTable

A foreign key

A stored procedure

Explanation:

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

A clustered index

A FileTable

A foreign key

A stored procedure

Reference:

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

Question: 76

HOTSPOT

To complete the sentence, select the appropriate option in the 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:

Explanation:

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: 77

HOTSPOT

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

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

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

Answer:

Explanation:

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

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

Reference:

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

Question: 78

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. (Choose three.)

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:

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

Question: 79

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

ProductName	Quantity
Product 1	100
Product?	129
Products	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:

Reference:

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

Question: 80

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: 81

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:

Reference:

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

Question: 82

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.

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

Answer:

Explanation:

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

Reference:

[https://en.wikipedia.org/wiki/Table_\(database\)#:~:text=A%20table%20is%20a%20collection,table%2](https://en.wikipedia.org/wiki/Table_(database)#:~:text=A%20table%20is%20a%20collection,table%2)

[Oformat%20within%20a%20database.&text=In%20relational%20databases%2C%20and%20flat,a%20row%20and%20column%20intersect.](https://en.wikipedia.org/wiki/View_(SQL))

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

https://en.wikipedia.org/wiki/Database_index#:~:text=A%20database%20index%20is%20a,maintain%20the%20index%20data%20structure.

Question: 83

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

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

Answer: A

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/az-cli-script-samples-content-guide?tabs=singledatabase>

Question: 84

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.

Statements

Yes

No

Azure SQL Database includes a fully managed backup service.

☐

☐

Azure SQL Database has built-in high availability.

☐

☐

Azure SQL Database can use Azure Advanced Threat Protection (ATP).

☐

☐

Answer:

Explanation:

Statements

Yes

No

Azure SQL Database includes a fully managed backup service.

☐

☐

Azure SQL Database has built-in high availability.

☐

☐

Azure SQL Database can use Azure Advanced Threat Protection (ATP).

☐

☐

Reference:

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

<https://azure.microsoft.com/en-us/blog/announcing-sql-atp-and-sql-vulnerability-assessment-general-availability/>

Question: 85

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.

Statements

Yes No

You can use Azure Data Studio to query a Microsoft SQL Q Server big data cluster.

☐

You can use Microsoft SQL Server Management Studio (SSMS) to query' an Azure Synapse Analytics data warehouse.

☐

☐

You can use MySQL Workbench to query' Azure Database for MariaDB databases.

☐

Answer:

Explanation:

Statements	Yes	No
------------	-----	----

You can use Azure Data Studio to query a Microsoft SQL Server big data cluster.		
---	--	--

You can use Microsoft SQL Server Management Studio (SSMS) to query an Azure Synapse Analytics data warehouse.		
---	--	--

You can use MySQL Workbench to query Azure Database		
---	--	--

for MariaDB databases.

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: 86

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.

Statements	Yes	No
------------	-----	----

You must apply patches to Azure SQL databases regularly.		
--	--	--

You need a Microsoft 365 subscription to create an Azure SQL database.		
--	--	--

You can use existing Microsoft SQL Server licenses to reduce the cost of Azure SQL databases.		
---	--	--

Explanation:

Answer:

Statements

Yes No

You must apply patches to Azure SQL databases regularly.

☐☒

You need a Microsoft 365 subscription to create an Azure SQL database.

☐☒

You can use existing Microsoft SQL Server licenses to reduce the cost of Azure SQL databases.

☒☐

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: 87

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

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

Answer: D

Explanation:

Reference:

<https://www.geeksforgeeks.org/difference-between-ddl-and-dml-in-dbms/>

Question: 88

HOTSPOT

To complete the sentence, select the appropriate option in the 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:

Explanation:

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 sender to protect data at rest.

Reference:

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

Question: 89

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:

Reference:

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

Question: 90

HOTSPOT

To complete the sentence, select the appropriate option in the 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:

Explanation:

By default, each Azure SQL database is protected by

a network security group (NSG), a server-level firewall.

Azure Firewall.

Azure Front Door.

Reference:

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

Question: 91

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:

Reference:

<https://docs.microsoft.com/en-us/sql/ssdt/project-oriented-offline-database-development?view=sql-serverver15>

Question: 92

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.

Components

Answer Area

Authentication

Firewall

Encryption

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:

Explanation:

Firewall

Prevent access to an Azure SQL database from another network.

Authentication
Encryption

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

Ensure that sensitive data never appear's as plain text in an Azure SQL database.

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: 93

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.

Statements

Yes No

Azure Table storage supports multiple read replicas.

0 0

Azure Table storage supports multiple write regions.

0 0

The Azure Cosmos DB Table API supports multiple read replicas.

o 0

The Azure Cosmos DB Table API supports multiple write regions.

0 0

Answer:

Explanation:

Statements

Yes

No

Azure Table storage supports multiple read replicas.

☐

☐

Azure Table storage supports multiple write regions.

☐

☐

The Azure Cosmos DB Table API supports multiple read replicas.

☐

☐

The Azure Cosmos DB Table API supports multiple write regions.

☐

☐

Reference:

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

Question: 94

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.

Data Store Types

Answer Area

Graph
Object
Key/value

Application users and their default language

Medical images and their associated metadata

Employee data that shows the relationships between employees

Answer:

Explanation:

Key/value
Object
Graph

Application users and their default language

Medical images and their associated metadata

Employee data that shows the relationships between employees

Reference:

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

Question: 95

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. (Choose two.)

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: 96

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:

Question: 97

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:

Reference:

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

Question: 98

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:

Question: 99

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:

Question: 100

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

Cassandra API
Gremlin API
MongoDB API
Table API

Answer Area

Graph data

JSON documents

Key/value data

Answer:

Explanation:

Gremlin API

Graph data

MongoDB API

JSON documents

Table API

Key/value data

Reference:

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

Question: 101

HOTSPOT

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

Answer:

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

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.

enable the hierarchical namespace.

set Replication kind to Read-access geo-redundant storage (RA-GRS).

set Performance to Premium.

set Replication to Read-access geo-redundant storage (RA-GRS).

Reference:

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

Question: 102

HOTSPOT

Answer:

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.

Reference:

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

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

Question: 103

You can query a graph database in Azure Cosmos DB

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

Explanation:

You can query a graph database in Azure Cosmos DB

- 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: 104

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:

Reference:

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

Question: 105

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.

Statements

Yes

No

When ingesting data from Azure Data Lake Storage across Azure regions, you will incur costs for bandwidth.

0

0

You can use blob, table, and file storage in the same Azure Storage account.

0

0

You implement Azure Data Lake Storage by creating an Azure Storase account.

0

0

Answer:

Explanation:

Statements

Yes ;

No

When ingesting data from Azure Data Lake Storage across Azure regions, you will incur costs for bandwidth.

1°

0

You can use blob, table, and file storage in the same Azure Storage account.

0 1 0

You implement Azure Data Lake Storage by creating an Azure Storase account.

0 1

0

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: 106

HOTSPOT

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

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



graph.

table.

partition key. document.

Answer:

Explanation:

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



- graph.
- table.
- partition key.
- document.

Reference:

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

Question: 107

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:

Question: 108

HOTSPOT

To complete the sentence, select the appropriate option in the 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:

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.

Question: 109

HOTSPOT

Answer:

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

NOTE: Each correct selection is worth one point.

Statements

Yes No

Processing salary' payments once a month is an example of a batch workload.

☐☐

A wind turbine that sends 50 sensor readings per second is an example of a streaming workload.

☐☐

A home electricity meter that sends readings once a day to an energy provider is an example of a streaming workload.

☐☐

Answer:

Explanation:

Statements	Yes	No
Processing salary payments once a month is an example of a batch workload.		
A wind turbine that sends 50 sensor readings per second is an example of a streaming workload.		
A home electricity meter that sends readings once a day to an energy provider is an example of a streaming workload.		

Question: 110

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:

Reference:

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

Question: 111

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

What should you do on the 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:

Reference:

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

Question: 112

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:

Reference:

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

Question: 113

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.

Activity Types

Control

Data movement

Data transformation

Answer Area


Copy

Mapping data flow

Until

Answer:

Explanation:



Reference:

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

Question: 114

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

NOTE: Each correct selection is worth one point.

- A. denormalized data
- B. heavy writes and moderate reads
- C. light writes and heavy reads
- D. schema on write
- E. schema on read
- F. normalized data

Answer: BDF

Explanation:

Reference:

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

Question: 115

Which two activities can be performed entirely by using the Microsoft Power BI service? Each correct answer presents a complete solution. (Choose two.)

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: A, D

Explanation:

Question: 116

HOTSPOT

To complete the sentence, select the appropriate option in the 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:

Explanation:

In Azure Data Factory; you can use

a control flow
a dataset
a linked sendee
an integration runtime

to orchestrate pipeline

activities that depend on the output of other pipeline activities

Reference:

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

Question: 117

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: A

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.

Reference:

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

Question: 118

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

NOTE: Each correct selection is worth one point.

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

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: 119

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.

Statements

Yes

No

A Microsoft Power BI dashboard is associated with a single workspace.

☐☐

A Microsoft Power BI dashboard can only display visualizations from a single dataset.

☐☐

A Microsoft Power BI dashboard can display visualizations from a Microsoft Excel workbook.

☐☐

Answer:

Explanation:

Statements

Yes ^o

A Microsoft Power BI dashboard is associated with a single workspace.



A Microsoft Power BI dashboard can only display visualizations from a single dataset.



A Microsoft Power BI dashboard can display visualizations from a Microsoft Excel workbook.



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: 120

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

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

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: 121

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

Analytics data warehouse.

What should you use to query the files?

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

Answer: C

Explanation:

Reference:

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

Question: 122

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.

Statements	Yes	No
Azure Databricks is an Apache Spark-based collaborative analytics platform.	0	0
Azure Analysis Services is used for transactional workloads.	0	0
Azure Data Factory orchestrates data integration workflows.	0	0

Explanation:

Answer:

Statements	Yes	No
Azure Databricks is an Apache Spark-based collaborative analytics platform.	0	0
Azure Analysis Services is used for transactional workloads.	0	0
Azure Data Factory orchestrates data integration workflows.	0	0

Question: 123

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.

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

Answer:

Explanation:

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

Box 1: Yes

Big data solutions often use long-running batch jobs to filter, aggregate, and otherwise prepare the data for analysis. Usually these jobs involve reading source files from scalable storage (like HDFS, Azure Data Lake Store, and Azure Storage), processing them, and writing the output to new files in scalable storage.

Box 2: No

Box 3: No

Reference:

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

Question: 124

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.

Visualizations

Answer Area

Key influencer

Scatter

Treemap

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:

Explanation:

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 values

A chart that shows the relationship between two numerical

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: 125

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.

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

Answer:

Explanation:

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

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: Yes

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>

<https://feedback.azure.com/forums/217321-sql-database/suggestions/6931152-please-add-ability-to-temporarily-turn-off-on-sql>

Development tools,
database management
business analytics

Networking
firewalls/security

Data center physical
plant/building

Question: 126

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.

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

Answer:

Explanation:

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

Reference:

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

Question: 127

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: 128

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.

Statements	Yes	No
The Azure Cosmos DB API is configured separately for each database in an Azure Cosmos DB account.	☐	☐
Partition keys are used in Azure Cosmos DB to optimize queries.	☐	☐
Items contained in the same Azure Cosmos DB logical partition can have different partition keys.	☐	☐

Answer:

Explanation:

Statements	Yes	No
The Azure Cosmos DB API is configured separately for each database in an Azure Cosmos DB account.	☐	☐
Partition keys are used in Azure Cosmos DB to optimize queries.	☒	☐
Items contained in the same Azure Cosmos DB logical partition can have different partition keys.	☐	☐

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: 129

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

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: 130

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.



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/sql-dw-walkthrough>

Question: 131

HOTSPOT

To complete the sentence, select the appropriate option in the 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:

Explanation:

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

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: 132

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.

Statements	Yes	No
A pipeline is a representation of a data structure within Azure Data Factory	☐	☐
Azure Data Factory pipelines can execute other pipelines	☐	☐
A processing step within an Azure Data Factory pipeline is an activity	☐	☐

Answer:

Explanation:

Statements	Yes	No
A pipeline is a representation of a data structure within Azure Data Factory	☐	☐
Azure Data Factory pipelines can execute other pipelines	☐	☐
A processing step within an Azure Data Factory pipeline is an activity	☐	☐

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: 133

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.

Services

Answer Area



Output data to Parquet format

Store data that is in Parquet format

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

Answer:

Explanation:

Azure Data Factory

Output data to Parquet format

Azure Data Lake Storage

Store data that is in Parquet format

Azure Synapse Analytics

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

Box 1: Azure Data Factory

Box 2: Azure Data Lake Storage

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

Box 3: Azure Synapse Analytics

Use Azure Synapse Analytics Workspaces.

Reference:

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

Question: 134

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.

Statements

Yes

Azure Synapse Analytics scales storage and compute independently	Q	Q
Azure Synapse Analytics can be paused to reduce compute costs	O	O
An Azure Synapse Analytics data warehouse has a fixed storage capacity	O	O

Answer:

Explanation:

Statements

Yes

Azure Synapse Analytics scales storage and compute independently	Q	Q
Azure Synapse Analytics can be paused to reduce compute costs	O	O
An Azure Synapse Analytics data warehouse has a fixed storage capacity	O	O

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: 135

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.

Statements	Yes	No
Azure Data Studio can be used to query an Azure SQL database from a device that runs macOS.	0	0
Microsoft SQL Server Management Studio (SSMS) enables users to create and use SQL notebooks.	0	0
Azure Data Studio can be used to restore a database.	0	0

Answer:

Explanation:

Statements	Yes	No
Azure Data Studio can be used to query an Azure SQL database from a device that runs macOS.	0	0
Microsoft SQL Server Management Studio (SSMS) enables users to create and use SQL notebooks.	0	0
Azure Data Studio can be used to restore a database.	0	0

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: 136

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.

Statements		Yes	No
Azure Databricks can consume	data from Azure SQL Database	☐	☐
Azure Databricks can consume	data from Azure Event Hubs	☐	☐
Azure Databricks can consume	data from Azure Cosmos DB	☐	☐

Answer:

Explanation:

Statements		Yes	No
Azure Databricks can consume data from Azure SQL Database		☐	☐
Azure Databricks can consume data from Azure Event Hubs		☐	☐
Azure Databricks can consume data from Azure Cosmos DB		☐	☐

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: 137

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.

Services

Azure Blob storage

Azure Cosmos DB

Azure Files

Azure Table storage

Answer Area

Service

Enables the use of SQL queries against data stored in JSON documents

Service

Enables users to access data by using the Server Message Block (SMB) version 3 protocol

Answer:

Explanation:

Azure Cosmos DB

Enables the use of SQL queries against data stored in JSON documents

Azure Files

Enables users to access data by using the Server Message Block (SMB) version 3 protocol

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: 138

Which component of Azure Data Factory enables you to define a sequence of activities including for each iterators?

- A. control flow
- B. linked services
- C. datasets
- D. the integration runtime

Answer: A

Explanation:

Question: 139

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

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

Answer: A

Explanation:

Reference:

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

Question: 140

Which three actions are used when performing diagnostic analytics?

A.

- Identify data anomalies
 - Collect anomaly-related data
 - Discover relationships and trends that explain anomalies.
- B.
- Identify data anomalies
 - Predict when anomalies are likely to recur
 - Use analytical tools such as decision trees and regression.
- C.
- Apply self-learning loops to draw inferences from anomalies.
 - Predict when anomalies are likely to recur
 - Use analytical tools such as decision trees and regression.
- D.
- Identify data anomalies
 - Use machine learning to find patterns in large datasets.
 - Discover relationships and trends that explain anomalies.

Answer: A

Explanation:

Question: 141

Which three actions are used when performing diagnostic analytics?

- A.
- Identify data anomalies
 - Collect anomaly-related data
 - Discover relationships and trends that explain anomalies.
- B.
- Identify data anomalies
 - Predict when anomalies are likely to recur
 - Use analytical tools such as decision trees and regression.
- C.
- Apply self-learning loops to draw inferences from anomalies.
 - Predict when anomalies are likely to recur
 - Use analytical tools such as decision trees and regression.
- D.
- Identify data anomalies
 - Use machine learning to find patterns in large datasets.
 - Discover relationships and trends that explain anomalies.

Answer: A

Explanation:

Question: 142

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

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

Answer: A

Explanation:

Azure Data Lake Storage Gen2 implements an access control model that supports both Azure rolebased 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: 143

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.

Locations

Answer Area

A standalone data analysis tool

Extract:

The CRM system

Load:

The data warehouse

Transform:

Answer:

Explanation:

Extract:

The CRM system

Load:

The data warehouse

Transform:

A standalone data analysis tool

Box 1: The CRM system


Data is extracted from the CRM system.

Box 2: The data warehouse

Data is loaded to the data warehouse.

Box 3: A standalone data analysis tool

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



Reference:

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

Question: 144

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

Reference:

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

Question: 145

HOTSPOT

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

Creating closed caption text for audio files
is an example of

	▼
cognitive	
descriptive	
predictive	
prescriptive	

analytics.

Answer:

Explanation:

Creating closed caption text for audio files
is an example of

	▼
cognitive	
descriptive	
predictive	
prescriptive	

analytics.

Descriptive, to answer the

question: What's

happening?

<7'0,2®

Descriptive Diagnostic
Predictive Prescriptive

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

Reference:

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

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

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

Question: 146 HOTSPOT

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

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

▼ analytics.

cognitive descriptive predictive
prescriptive

Answer:

Explanation:

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

▼ analytics.

cognitive descriptive predictive


prescriptive

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?



Descriptive Diagnostic Predictive Prescriptive

What's happening? Why's happening? What will happen? 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: 147 HOTSPOT

To complete the sentence, select the appropriate option in the 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:

Explanation:

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.

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: 148

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

Which command should you run?

- A. az scorage share create
- B. az scorage account creace
- C. az cosmosdb creace
- D. az scorage concainer creace

Answer: D

Explanation:

Reference:

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

Question: 149

HOTSPOT

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

The Azure Cosmos DB ▼ API enables the use of SELECT statements to Core (SQL) Gremlin MongoDB Table
retrieve documents from Azure Cosmos DB.

Answer:

Explanation:

The Azure Cosmos DB _____ ▼ API enables the use of SELECT statements to
Core (SQL)
Gremlin
MongoDB
Table

retrieve documents from Azure Cosmos DB.

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: 150

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:

Question: 151

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.

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

Answer:

Explanation:



Reference:

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

Question: 152

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.

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:

Explanation:

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)

Reference:

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

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

Question: 153

HOTSPOT

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



Column family databases

natively support the analysis of relationships between entities

Document databases
Graph databases
Key-value stores

Answer:

Explanation:

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

natively support the analysis of relationships between entities

Question: 154
HOTSPOT

To complete the sentence, select the appropriate option in the 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:

Explanation:

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: 155
HOTSPOT

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

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

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

Answer:

Explanation:

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

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

Reference:

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

Question: 156
HOTSPOT

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

A block of code that runs in a database is called

a stored procedure.

a table.

a view.

an index.

Answer:

Explanation:

A block of code that runs in a database is called

a stored procedure a table.

a view.

an index.

Reference:

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

Question: 157

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

Reference:

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

Question: 158

HOTSPOT

To complete the sentence, select the appropriate option in the 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:

Explanation:

By default, each Azure SQL database is protected by

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

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: 159

HOTSPOT

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

When provisioning an Azure Cosmos DB . you need to specify which type of API you will use
account container database item

Answer:

Explanation:

When provisioning an Azure Cosmos DB y, you need to specify which type of API you will use
account container database item

Reference:

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

Question: 160

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: 161

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: 162

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.

Statements

Yes

No

Azure Table storage within a single Azure Storage account supports multiple concurrent reads in different Azure regions

0

0

Azure Table storage within a single Azure Storage account supports multiple concurrent writes in different Azure regions

0

^

An Azure Cosmos DB account that uses the Table API supports multiple concurrent reads in different Azure regions

~0~

[0

An Azure Cosmos DB account that uses the Table API supports multiple concurrent writes in different Azure regions.

0

0

Answer:

Explanation:

Statements

yes No

Azure Table storage within a single Azure Storage account supports multiple concurrent reads in different Azure regions.

☐

☒

Azure Table storage within a single Azure Storage account supports multiple concurrent writes in different Azure regions.

☐

An Azure Cosmos DB account that uses the Table API supports multiple concurrent reads in different Azure regions.

☐

An Azure Cosmos DB account that uses the Table API supports multiple concurrent writes in different Azure regions.

☐

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: 163

HOTSPOT

To complete the sentence, select the appropriate option in the 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:

Explanation:

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).

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: 164

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: 165

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.

Statements	Yes	No
You can copy a dashboard between Microsoft Power BI workspaces.	☐	☐
A Microsoft Power BI dashboard can only display visualizations from a single dataset	☐	☐
A Microsoft Power BI dashboard can display visualizations from a Microsoft Excel workbook	☐	☐

Answer:

Explanation:

Statements	Yes	No
You can copy a dashboard between Microsoft Power BI workspaces.	☐	☐
A Microsoft Power BI dashboard can only display visualizations from a single dataset	☐	☐
A Microsoft Power BI dashboard can display visualizations from a Microsoft Excel workbook	☐	☐

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: 166

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: C

Explanation:

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

Query both relational and nonrelational data.

Reference:

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

Question: 167

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.

Statements	Yes	No
Azure Databricks is an Apache Spark-based collaborative . analytics platform.	☐	☐
Azure Analysis Sendees is used for transactional workloads.	☐	☐
Azure Data Factory' orchestrates data integration workflows.	☐	☐

Answer:

Explanation:

Statements

Yes

No

Azure Databricks is an Apache Spark-based collaborative analytics platform.

☐☐

Azure Analysis Sendees is used for transactional workloads.

☐☒

Azure Data Factory orchestrates data integration workflows.

☐☐

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: 168

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.

Analytics Types

Answer Area

Cognitive

Diagnostic

Descriptive

Predictive

Prescriptive

Why did sales increase last month?

Provide specific guidance on how to allocate current budget to buy different inventory items.

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

Answer:

Explanation:

Box 1: Diagnostic

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

Box 2: Prescriptive

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

Box 3: Cognitive

Reference:

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

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

Question: 169

HOTSPOT

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

▼ presents content defined by a query

A heap

A stored procedure

A view

An index

Answer:

Explanation:

A View

Question: 170

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:

Reference:

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

Question: 171

HOTSPOT

Select the answer that correctly completes the sentence.

A ▼ is a chart of colored rectangles. The size of the rectangles represent line chart matrix scatter treemap

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

Answer:

Explanation:

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

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: 172

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:

Question: 173

HOTSPOT

Select the answer that correctly completes the sentence.

In a SQL statement,

▼ are examples of aggregate functions

FROM and WHERE

GROUP BY and ORDER BY

JOIN and MERGE

MAX and MIN

Answer:

Explanation:

In a SQL statement,

▼ are examples of aggregate functions.

FROM and WHERE

GROUP BY and ORDER BY

JOIN and MERGE

MAX and MIN

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

GROUPINGID

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: 174
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.

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:

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: 175

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.

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

Answer:

Explanation:

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

Reference:

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

Question: 176

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:

Reference:

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

Question: 177

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

Sales by Year and Quarter

- Increase • Decrease • Total

2013Qtr3 2013QU4 2014Qtr 1 2014QU2 2014Qtr3 2014QU4 Total



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: 178

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

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

Answer: D

Explanation:

Question: 179

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: C

Explanation:

Question: 180

What is a characteristic of a non-relational database?

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

Answer: B

Explanation:

Question: 181

Select the answer that correctly completes the sentence.

Answer Area

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

Answer: See the
answer in
explanation below.

Explanation:

Answer below as:

Answer Area

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

Question: 182

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,

Answer: See the
answer in
explanation below.

Explanation:

Answer below as:

Answer Area

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

Question: 183

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.

Azure SQL services

Answer Area

Azure SQL Database elastic pool

Azure SQL Database single database

Azure SQL Managed Instance

A serverless configuration

A database that has the highest compatibility with on-premises Microsoft SQL Server

A low-cost low-maintenance database

Answer:

Explanation:

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

Question: 184

Select the answer the correctly completes the sentence.

Answer Area

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

Answer: See the answer in explanation below.

Explanation:

Answer below as:

Answer Area

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

Question: 185

Select the answer that correctly completes the sentence.

Answer Area

☐ A block of code that runs in a database is called

Answer: See the answer in explanation below.

Explanation:

Answer below as:

Answer Area

☐ A block of code that runs in a database is called a stored procedure. ☐

Question: 186

Select the answer that correctly completes the sentence.

Answer Area

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

Answer: See the
answer in
explanation below.

Explanation:

Answer below as:

Answer Area

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

Question: 187

You have data saved in the following format:

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

Which format was used?

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

Answer: A

Explanation:

Question: 188

What is a primary characteristic of a relational database?

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

Answer: A

Explanation:

Question: 189

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

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

Answer: C

Explanation:

Question: 190

Select the answer that correctly completes the sentence.

Answer Area

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

Answer: See the
answer in
explanation below.

Explanation:

Answer below as:

Answer Area

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

Question: 191

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

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

Answer: D

Explanation:

Question: 192

Select the answer that correctly completes the sentence.

Answer Area

A JSON document is an example of

Answer: See the
answer in
explanation below.

Explanation:

Answer below as:

Answer Area

A JSON document is an example of semi-structured data.

Question: 193

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

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

Answer: B

Explanation:

Question: 194

Select the answer that correctly completes the sentence.

Answer Area

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

Answer: See the
answer in
explanation below.

Explanation:

Answer below as:

Answer Area

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

Question: 195

Select the answer that correctly completes the sentence.

Answer Area

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

Answer: See the answer in explanation below.

Explanation:

Answer below as:

Answer Area

An instance of SQL Server on Azure Virtual Machines ^ can be used as part of a lift-and-shift solution for migrating an on-premises application to the cloud with minimal changes.

Question: 196

Select the answer that correctly completes the sentence.

Answer Area

Microsoft SQL Server database tables store

Answer: See the answer in explanation below.

Explanation:

Answer below as:

Answer Area

Microsoft SQL Server database tables store relational data.

Question: 197

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

- A. XML
- B. KQL

- C. YAML
- D. SQL

Answer: D

Explanation:

Question: 198

Which type of Azure resource supports the serverless configuration of an Azure SQL database?

- A. Azure SQL Managed Instance
- B. SQL Server on Azure Virtual Machines
- C. a single database in Azure SQL Database
- D. an Azure SQL Database elastic pool

Answer: B

Explanation:

Question: 199

Select the answer that correctly completes the sentence.

Answer: See the
answer below:

Explanation:

Answer as below.

Answer Area

Backing up and restoring data are the responsibilities of a database administrator

Question: 200

Select the answer that correctly completes the sentence.

Answer: See the
answer below:

Explanation:

Answer as below.

Answer Area

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

that hosts SQL Server.

Question: 201

correctly completes the sentence.

Answer: See the answer below:

Explanation:

Answer as below.

Answer Area

An elastic pool

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

Question: 202

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

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

Answer: B

Explanation:

Question: 203

You have data saved in the following format: (Person Agt*23">

(Fir Ft Name>?ohn(/FirstName>

(LastName>S<ith(/LastName>

(Hobbies*

(Hobby Type="Sports">Golf</Hobby>

(Hobby Type="LEisure">Reading</Hobby*

(Hobby Type*"%eisure%Mk4tar(/Hobby>

</Hobbies* (/Parsons

Which format was used?

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

Answer: B

Explanation:

Question: 204

Select the answer that correctly completes the sentence.

Answer: See the
answer below:

Explanation:

Answer as below.

Answer Area

In a graph database, edges specify relationships between tables

Question: 205

Select the answer that correctly completes the sentence.

Answer: See the
answer below:

Explanation:

Answer as below.

Answer Area

In a data warehousing workload data from multiple sources is combined in a single location

Question: 206

Select the answer that correctly completes the sentence.

Answer: See the
answer below:

Explanation:

Answer as below.

Answer Area

In batch processing delay in delivering data processing results is acceptable⁷

Question: 207

Select the answer that correctly completes the sentence.

Answer: See the
answer below:

Explanation:

Answer as below.

Answer Area

A key/value data store is optimized for simple lookups.

Question: 208

Select the answer that correctly completes the sentence.

Answer: See the
answer below:

Explanation:

Answer as below.

Answer Area

An Azure Storage account must have Hierarchical namespace * enabled to support Azure Data Lake Storage

Question: 209

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

Answer: A

Explanation:

Question: 210

Select the answer that correctly completes the sentence.

Answer: See the
answer below:

Explanation:

Answer as below.

Answer Area

You can query a graph database in Azure Cosmos DB as nodes and edges by using the Gremlin language

Question: 211

Select the answer that correctly completes the sentence.

Answer: See the
answer below:

Explanation:

Answer as below.

Answer Area

A non-relational database stores data without a fixed schema

Question: 212

Select the answer that correctly completes the sentence.

Answer: See the
answer below:

Explanation:

Answer as below.

Answer Area

A relational database must be used when strong consistency guarantees are required.

Question: 213

Select the answer that correctly completes the sentence.

Answer: See the
answer below:

Explanation:

Answer as below.

Answer Area

Relational data is stored in a tabular form of rows and columns.

*

Question: 214

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

Answer: A, C, F

Explanation:

Question: 215

Select the answer that correctly completes the sentence.

Answer: See the
answer below:

Explanation:

Answer as below.

Answer Area

In a relational database, each row in a table has a key-value pair.

Question: 216

Select the answer that correctly completes the sentence.

Answer: See the
answer below:

Explanation:

Answer as below.

Answer Area

A clustered index ' is an object associated with a table that sorts
and stores data rows based on their key values.

Question: 217

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

An Azure Storage account must have

Blob soft delete
Hierarchical namespace
Large file shares
Versioning

enabled to support Azure Data Lake Storage.

Answer:

Explanation:

Hierarchical namespace

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

Question: 218

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. isolation
- C. durability
- D. consistency

Answer: B

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: 219

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 Edge
- D. Azure SQL Database

Explanation:

Answer: A

Question: 220

A data engineer is responsible for which task?

- A. Design and build analytical models.
- B. Design and implement data stores for analytical workloads.
- C. Explore data to identify trends.
- D. Implement policies, tools, and processes for backup and recovery plans.

Answer: D

Explanation:

Question: 221

HOTSPOT

Select the answer that correctly completes the sentence.

Select the answer that correctly completes the sentence.

Answer Area

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

Microsoft Power BI Desktop

The Microsoft Power BI service The Microsoft Power BI phone app

Answer:

Explanation:

Answer Area

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

Question: 222

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

Normalizing a database increases the throughput of writing transactions.

Analytics systems are more normalized than transactional systems.

Normalizing a database results in queries that require more joins.

Answer:

Explanation:

YES, NO, YES Normalizing decreases throughput for reading (need to make joins), but increases for writing (Less to write, since there are no duplicates)

Question: 223

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

An extract, transform, and load (ETL) process requires data that is fully processed before being loaded to the target data store. A matching schema in the data source and the data target is critical for the data store powerful enough to transform data. _____
that the data target be a relational database. _____
data that is fully processed before being loaded to the target data store. _____

Answer:

Explanation:

Answer Area

An extract, transform, and load (ETL) process requires data that is fully processed before being loaded to the target data store.

Question: 224
HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

In Microsoft Power BI, a

dashboard

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

dashboard

dataset report tile

Answer:

Explanation:

Answer Area

In Microsoft Power BI, a dashboard is a collection of visualizations that appear together on one or more pages.

Question: 228
HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area		
In a relational database, the	primary key *	uniquely identifies each row in a table.
	foreign key	
	primary key	
	WHERE clause	
	nonclustered index	

Answer:

Explanation:

Answer Area		
In a relational database, the	primary key	uniquely identifies each row in a table.

Question: 229
HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area	
	presents content defined by a query.
A view A heap A stored procedure A view An index	

Explanation:

Answer Area	
A view	presents content defined by a query.

Answer:

Question: 230
HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

natively support the analysis of relationships between entities.

Answer:

Explanation:

Answer Area

Column family databases

natively support the analysis of relationships between entities.

Question: 231

What is a characteristic of non-relational data?

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

Answer: D

Explanation:

Question: 232

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

A key/value data store is optimized for

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

Answer:

Explanation:

Answer Area

A key/value data store is optimized for

- simple lookups.

Question: 233

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. WHERE
- B. SELECT
- C. JOIN

Answer: B

Explanation:

Question: 234

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

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

Answer:

Explanation:

Answer Area

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

- optimized for transactions.

Question: 235

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

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

Answer:

Explanation:

Answer Area

Microsoft Power BI Desktop

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

Question: 236

Which type of database contains nodes and edges?

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

Answer: A

Explanation:

Question: 237

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

- A data engineer
- 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:

Explanation:

Answer Area


- A data engineer

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

Question: 238


HOTSPOT

Select the answer that correctly completes the sentence.



Answer:

Explanation:




Question: 239

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

NOTE: Each correct match is worth one point.



Answer:

Explanation:



Question: 240

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

Relational data is stored in a tabular form of rows and columns, a file system as unstructured data, a hierarchal folder structure.

a tabular form of rows and columns.

I comma-separated value (CSV) files.

Answer:

Explanation:

Answer Area

Relational data is stored in a tabular form of rows and columns.

Question: 241

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

A relational database must be used when

strong consistency quarantees are required.

a dynamic schema is required.

data will be stored as key/value pairs.

storing large images and videos.

strong consistency quarantees are required.

Answer:

Explanation:

Answer Area

A relational database must be used when

strong consistency quarantees are required.

Question: 242

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

JavaScript
KQL
Python
SQL

Answer Area

Azure Data Explorer

Azure Stream Analytics

Azure Synapse Analytics

Explanation:

Language

JavaScript

Answer Area

Python

Azure Data Explorer

Python | KQL

SQL

SQL

Azure Stream Analytics

Azure Synapse Analytics

Question: 243 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

Ye*

No

You create Azure File Storage in a storage account.

The Premium tier of Azure File Storage uses HDD-based hardware in a datacenter.

The Standard tier of Azure File Storage uses SSDs.

Answer:

Explanation:

Answer Area

Statements

Yes

No

You create Azure File Storage in a storage account.

The Premium tier of Azure File Storage uses HDD-based hardware in a datacenter.

The Standard tier of Azure File Storage uses SSDs.

Question: 244 HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

Microsoft Power BI Desktop

Microsoft Power BI Desktop

The Microsoft Power BI service

The Microsoft Power BI phone app

* provides a fun range of Power BI data modeling and report editing feature*

Answer:

Explanation:

Question: 245

You have the following two tables of customer data.

CustomcrID	NameStylc	Title	FirstName	MiddlcNamc	Last Na me
29486		0Ms	Kim	NULL	Abercrombie
29489		0Ms.	Frances	e	Adams
29490		0Ms	Margaret	J	Smith
29492		0Mi	Jay	NULL	Adams

CustomcrID	AddressID	AddressType	ModificdDadc
29486	621	Main Office	9/1/05 0:00
29489	1069	Main Office	7/1/05 DOO
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. non-relational
- B. structured
- C. semi-structured
- D. unstructured

Answer: D

Explanation:

Question: 246

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

[An instance of SQL Server on Azure Virtual Machines] can be used as part of a lift-and-shift solution for migrating [An Azure SQL Database elastic pool] ranges
An instance of SQL Server on Azure Virtual Machines
A single database in Azure SQL Database

Answer:

Explanation:

Answer Area

An instance of SQL Server on Azure Virtual Machines can be used as part of a lift-and-shift solution for migrating an on-premises application to the cloud with minimal changes.

Question: 247

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

Answer:

Explanation:

Answer Area

A data analyst

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

Question: 248

Which file format has defined names and data types for each column and uses compressed columnar storage?

- A. JSON
- B. CSV
- C. Apache Parquet
- D. HTML

Answer: C

Explanation:

Question: 249

You have structured data in tabular format.

What represents an individual instance of a data entity?

- A. a column
- B. a cell
- C. a row
- D. a table

Answer: C

Explanation:

Question: 250

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 timestamp
- B. a range
- C. a row key
- D. a partition key

Answer: C, D

Explanation:

Question: 251

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:

Question: 252

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:

Explanation:

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

Question: 253

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.

- JSON
- CSV
- ORC
- TSV

Answer:

Explanation:

JSON

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

Question: 254

Which language is used to define queries in Azure Synapse Data Explorer?

- A. Bash
- B. PowerShell
- C. KQL
- D. SQL

Answer: C

Explanation:

Question: 255

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

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

Answer: B

Explanation:

Question: 256

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
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:

Explanation:

Answer Area

Statements

Database administrators apply data cleansing routines and turn data into useful information.

Yes

No

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.

Question: 257

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.

Processes	Answer Area
Data ingestion	Data is retrieved and processed before being saved to a target system.
Extract, load, and transform (ELT) processing	Data is saved to a target system, and then processed.
Extract, transform, and load (ETL) processing	Data is extracted from a source system, and then saved to a target system.

Answer:

Explanation:

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

J target system.

Question: 258

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area	Answer:
Explanation: A database administrator is responsible for managing the security of the data in a database. A data analyst A data engineer A data scientist A database administrator is responsible for managing the security of the data in a database. A database administrator	

Question: 259

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:

Question: 260

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



Which type of data store is this?

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

Answer: B

Explanation:

Question: 261

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

UPDATE ' is an example of a Data Manipulation Language (DML) statement. CREATE ,
DROP
RENAME
UPDATE

Answer:

Explanation:

Answer Area

UPDATE ' is an example of a Data Manipulation Language (DML) statement.

Question: 262

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

Database administrators ' are responsible for the availability, performance, and optimization of databases
Data analysts
Data engineers
Data scientists
Database administrators

Answer:

Explanation:

Answer Area

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

Question: 263

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:

Term
Data warehouses
Dimension tables
Extract, transform, and load (ETL)
Fact tables
Online analytical processing (OLAP)

Answer Area

Question: 264

DRAG DROP

Match the terms to the appropriate descriptions.

Explanation:

Term	Answer Area
Data warehouses	Data warehouses A solution that copies data from files and online transaction processing (OLTP) databases into an analytical processing system.
Dimension tables	Dimension tables Contains numeric values that can be analyzed.
Extract, transform, and load (ETL)	

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.

A solution that copies data from files and online transaction processing (OLTP) databases into an analytical processing system.
Contains numeric values that can be analyzed.

A solution that stores data in a relational schema that is optimized for read operations

Fact tables

Fact tables
A solution that stores data in a relational schema that is optimized for read operations.
Online analytical processing (OLAP)

Answer:

Question: 265
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 a transactional database that guarantees atomicity, all statements in a transaction must complete successfully or the transaction must roll back.		
In a transactional database that guarantees consistency, all data changes will be logged and accounted for.		
In a transactional database that guarantees isolation, processes can see only data in a consistent state.		

Answer:

Explanation:

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.	☐	☒
In a transactional database that guarantees consistency, all data changes will be logged and accounted for.	☒	☐
In a transactional database that guarantees isolation, processes can see only data in a consistent state.	☒	☐

Question: 266

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

Database backup and recovery

Enabling advanced analytic capabilities through reports

Monitoring data pipelines

Answer:

Explanation:

Job role

Data analyst

Data engineer

Database administrator

Answer Area

Database administrator Database backup and recovery

Data engineer Enabling advanced analytic capabilities through reports

Data analyst Monitoring data pipelines

Question: 267

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 a relational database, the schema is defined before data is written to the table.	☐	☐
Key/value data stores are highly optimized for complex queries that contain multiple joins.	☐	☐
A relational database is commonly used to store and query structured data,	☐	☐

Explanation:
Answer Area

Answer:

Statements

- In a relational database, the schema is defined before data is written to the table.
- Key/value data stores are highly optimized for complex queries that contain multiple joins.
- A relational database is commonly used to store and query structured data.

Yes No

Question: 268

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

In a graph database, edges specify relationships between

tables.
column families.
facts.
vertexes.
tables.

Answer:

Explanation:

Answer Area

In a graph database, edges specify relationships between tables.

Question: 269

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

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

Answer:

Explanation:

Answer Area

The pipeline engine used in Azure Data Factory is also available in Azure Synapse Analytics.

Question: 270

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

The Azure Cosmos DB for NoSQL

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

lot April In* Gremlin for MongoDB for Table

Answer

Explanation:

Answer Area

The Azure Cosmos DB for NoSQL

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

Question: 271

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

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

Answer: A

Explanation:

Question: 272

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

A Microsoft Power BI paginated report^{T1} enables users to create highly formatted, dashboard interactive report

paginated report

| subscription | fixed-layout documents optimized for printing

and archiving.

Answer

Explanation:

Answer Area

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

Question: 273
HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

Batch workloads

collect data and then process the data when a condition is met.

always process data in memory, row-by-row.

collect and process data at most once a day.

process data as new data is received in near real-time.

collect data and then process the data when a condition is met.

Answer

Explanation:

Answer Area

Batch workloads

collect data and then process the data when a condition is met.

Question: 274
HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

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

a target data store powerful enough to transform data

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

Answer

Explanation:

Answer Area

An extract, load, and transform (ELT) process requires a target data store powerful enough to transform data.

Question: 275

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

Relational data uses

keys ^ | to enforce relationships between different tables,
collections
rows
partitions

Answer:

Explanation:

Answer Area

Relational data uses

keys

to enforce relationships between different tables.

Question: 276

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

semi-structured data.

semi-structured data.
structured data.
unstructured data,

Answer:

Explanation:

Answer Area

Varying fields for each entity in a JSON document is an example of semi-structured data.

Question: 277

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

semi-structured data.
relational

semi-structured

structured
unstructured

Answer:

Explanation:

Answer Area

A JSON document that contains customer entities with a variable number of properties is an example of semi-structured
r data.

Question: 278

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

In an analytical model, any point where the dimensions intersect represents
I a cube,
a cube
an aggregated dimension, an
aggregated measure.

Answer:

Explanation:

Answer Area

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

Question: 279

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

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

Answer:

Explanation:

Answer Area

A data analyst

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

Question: 280

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: B

Explanation:

Question: 281

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.

In non-relational databases, data is always stored by using a fated schema, in non-relational databases, two entities in the same collection can have different sets of fields.

Answer:

Explanation:

Answer Area

Statements

Yes No

In non-relational databases, data is always stored in tables

In non-relational databases, data is always stored by using a fixed schema.

In non-relational databases, two entities in the same collection can have different sets of fields.

Question: 282

You have a SQL query that combines customer data and order data.

a. The query includes calculated columns.

You need to create a database object that would allow other users to rerun the same SQL query. What should you create?

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

Answer: D

Explanation:

Question: 283

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

Python
Scala
SQL

Answer

Explanation:

Answer Area

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

Question: 284

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:

Question: 285

You have a banking application that transfers money in to and out of accounts.
Of which type of solution is this an example?

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

Answer: D

Explanation:

Question: 286

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

The Azure SQL Edge service is optimized for IoT solutions.

Azure SQL Database

Azure SQL Edge

Azure SQL Managed Instance

SQL Server on Azure Virtual Machines

Answer:

Explanation:

Answer Area

The Azure SQL Edge service is optimized for IoT solutions.

Question: 287

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

combines a scalable high performance Microsoft SQL Server based relational data warehouse

Azure ML

Azure HDInsight
Azure Stream Analytics
Azure Synapse Analytics

with the flexibility of a data lake and open-source Apache Spark.

Answer:

Explanation:

Answer Area

Azure Data Lake Storage combines the flexibility of a data lake and open-source Apache Spark with the flexibility of a data lake and open-source Apache Spark

Question: 288

HOTSPOT

Select the answer that correctly completes the sentence.

Azure Blob Storage page blobs are optimized for

random read and write operations.

storing key/value data
storing large binary objects that change infrequently

Answer:

Explanation:

Answer Area

Azure Blob Storage page blobs are optimized for

random read and write operations.

Question: 289

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

Atomicity

Consistency

Isolation

Durability

Concurrent transactions cannot interfere with one another and must result in a consistent database state.

Each transaction is treated as a single unit that succeeds or fails completely.

Completely consistent

Transactions can only take the data in a database from one valid state to another.

When a transaction is committed, it will remain committed

Answer:

Explanation:

Terms

Atomicity

Consistency

Isolation

Durability

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: 290

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

Create a pipeline to transfer data from a database to a data warehouse.

Create a transactional database.

Create a chart that shows sales transactions by week.

Answer:

Explanation:

Job roles

Data analyst

Data architect

Data engineer

Database administrator

Answer Area

Data engineer>Create a pipeline to transfer data from a database to a data warehouse.

Data architect>Create a transactional database.

Data analyst>Create a chart that shows sales transactions by week.

Question: 291

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:**Explanation:****Terms**

Attributes

Entities

Foreign keys

Primary keys

Schemas

Answer Area

Attributes	Get separated into individual columns
Entities	Get separated into individual tables
Foreign keys	Link two related entities

Question: 292

Which Azure Storage service implements the key/value model?

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

Answer: B**Explanation:****Question: 293**

How is raw data stored in a data lakehouse?

- A. as tables
- B. as queues
- C. as files
- D. as disks

Answer: B

Question: 294

Blob rehydration occurs in Azure Blob storage when a blob moves between which access tiers?

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

Answer: A

Question: 295

What is the typical latency for read operations against media files stored in the Cool access tier of Azure Blob Storage?

- A. a few milliseconds
- B. a few seconds
- C. a few minutes
- D. a few hours

Answer: A

Question: 296

Which three services support Apache Spark? Each correct answer presents a complete solution NOTE: Each correct selection is worth one point.

- A. Azure Data Explorer
- B. Azure Synapse Analytics
- C. Azure Databricks
- D. Azure Stream Analytics
- E. Azure HDInsight

Answer: A, B, C

Question: 297

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

A clustered index '1 physically sorts the data in a table based on the values in a specified column

A view

A clustered index

A stored procedure

A nonclustered index

Answer:

Answer Area

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

Question: 298

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

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

Of which type of data is this an example?

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

Answer: D

Question: 299

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

Data Store Types

Answer Area

.. Graph
./, key/value .
.. Object

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

Answer:

Data Store Types

Graph
Key/value
Object

Answer Area

Key/value	A list of application user names and, for each, the
Object	A collection of medical images and their associated metadata
Graph	Employee data that shows the relationships between

Question: 300

What can you use to create data ingestion pipelines?

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

Answer: C

Question: 301

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

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

Answer: D

Question: 302

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

Apache Spark Structured Streaming is built on a that is used to encapsulate a single table of data.



- delta lake
- data lake
- i data warehouse

Answer:

Answer Area

Apache Spark Structured Streaming is built on a dataframe ^ that is used to encapsulate a single table of data.

Question: 303

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.	☐	☐
Elastic pools are used to share resources across multiple instances of SQL Server on Azure Virtual Machines.	☐	☐
Azure SQL Database is fully compatible with both on-premises physical instances and virtualized instances of Microsoft SQL Server	☐	☐

Answer:

Answer Area

Statements	Yes	No
Each Azure SQL managed instance supports multiple databases that can be accessed by using cross-database queries.	☒	☐
Elastic pools are used to share resources across multiple instances of SQL Server on Azure Virtual Machines.	☒	☐
Azure SQL Database is fully compatible with both on-premises physical instances and virtualized instances of Microsoft SQL Server	☒	☐

Question: 304

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
Stream processing can process all the data in a dataset at once.	☐	☐
Batch processing occurs immediately, with latency in the order of seconds or milliseconds.	☐	☐
Stream processing is used for simple response functions, aggregates, or calculations, such as rolling averages.	☐	☐

Answer:

Answer Area

Statements

Yes

No

Stream processing can process all the data in a dataset at once.

☐

☐

Batch processing occurs immediately, with latency in the order of seconds or milliseconds.

☐

☐

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

☐

Question: 305

What is the responsibility of a database administrator?

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

Answer: C

Question: 306

You have data saved in the following format.

firsttname, LastName, Age, LeisureMobility, Sports Mote*, John, Smith, 23, Reading, Basketball
Ben, Saith, 21, Guitar, Curling

Which format was used?

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

Answer: D