

Exam A

QUESTION 1

HOTSPOT

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

Hot Area:

Answer Area

Descriptive analytics tells you

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

Correct Answer:

Answer Area

Descriptive analytics tells you

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

Section: Describe core data concepts

Explanation

Explanation/Reference:

Reference:

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

QUESTION 2

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
Normalization involves eliminating relationships between database tables.	<input type="radio"/>	<input type="radio"/>
Normalizing a database reduces data redundancy.	<input type="radio"/>	<input type="radio"/>
Normalization improves data integrity.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

Statements	Yes	No
Normalization involves eliminating relationships between database tables.	<input type="radio"/>	<input checked="" type="radio"/>
Normalizing a database reduces data redundancy.	<input checked="" type="radio"/>	<input type="radio"/>
Normalization improves data integrity.	<input checked="" type="radio"/>	<input type="radio"/>

Section: Describe core data concepts

Explanation

Explanation/Reference:

Reference:

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

QUESTION 3

HOTSPOT

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

Hot Area:

Answer Area

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.	

Correct Answer:

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.	

Section: Describe core data concepts

Explanation

Explanation/Reference:

Reference:

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

QUESTION 5

HOTSPOT

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

Hot Area:

Answer Area

Transcribing audio files is an example of

	▼
cognitive	
descriptive	
predictive	
prescriptive	

analytics.

Correct Answer:

Answer Area

Transcribing audio files is an example of analytics.

- cognitive
- descriptive
- predictive
- prescriptive

Section: Describe core data concepts

Explanation

Explanation/Reference:

Reference:

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

QUESTION 6

DRAG DROP

Match the types of analytics that can be used to answer the business questions.

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

NOTE: Each correct match is worth one point.

Select and Place:

Analytics Types

- Cognitive
- Diagnostic
- Descriptive
- Predictive
- Prescriptive

Answer Area

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

Correct Answer:

Analytics Types

- Cognitive
- Diagnostic
- Descriptive
- Predictive
- Prescriptive

Answer Area

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

Explanation/Reference:

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.

Incorrect Answer:

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

Box 3: Descriptive

Generally speaking, data analytics comes in four types:

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

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

Predictive, to answer the question: What will happen?

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

**Reference:**

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

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

QUESTION 7**HOTSPOT**

You have the following JSON document.


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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Explanation

DataWolfs

Explanation/Reference:

Reference:

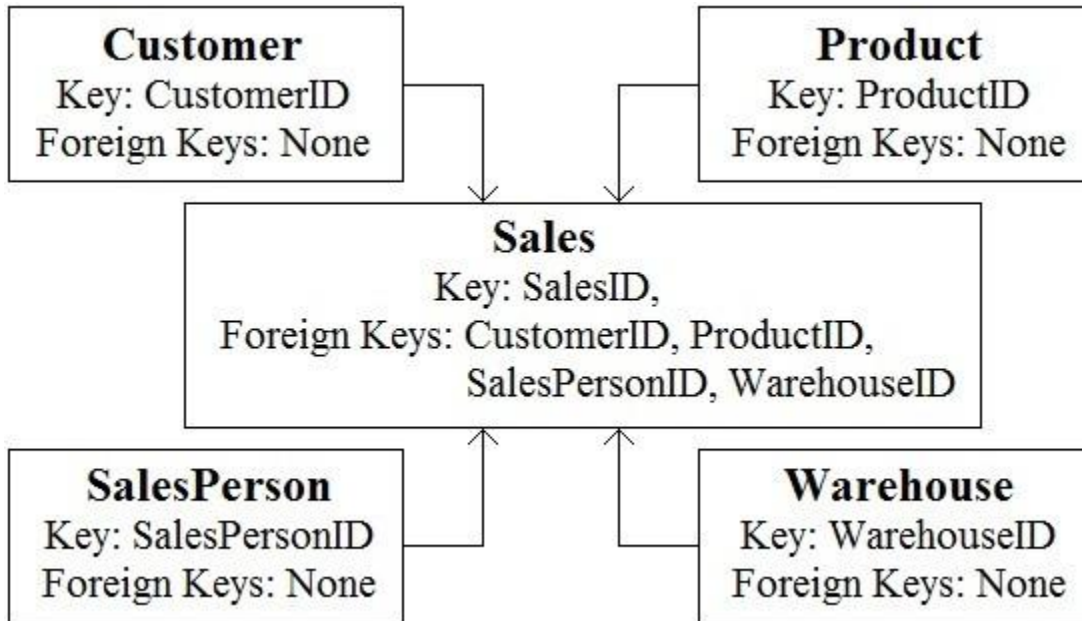
https://www.w3schools.com/js/js_json_arrays.asp

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

QUESTION 8

HOTSPOT

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



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

NOTE: Each correct selection is worth one point

Hot Area:

Answer Area

The data model is a [answer choice].

▼
transactional model
star schema
snowflake schema

Customer is a [answer choice] table.

▼
fact
dimension
bridge

Correct Answer:

Answer Area

The data model is a [answer choice].

	▼
transactional model	
star schema	
snowflake schema	

Customer is a [answer choice] table.

	▼
fact	
dimension	
bridge	

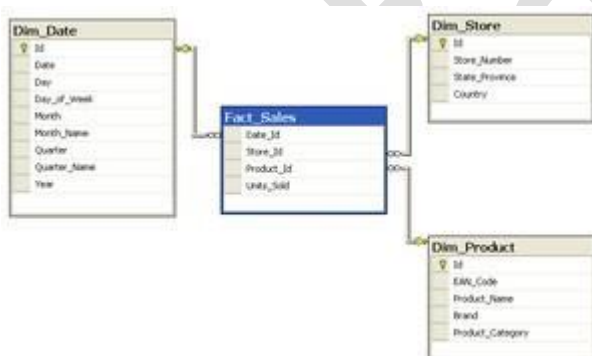
Section: Describe core data concepts
Explanation

Explanation/Reference:
 Explanation:

Box 1: star schema

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

Example:

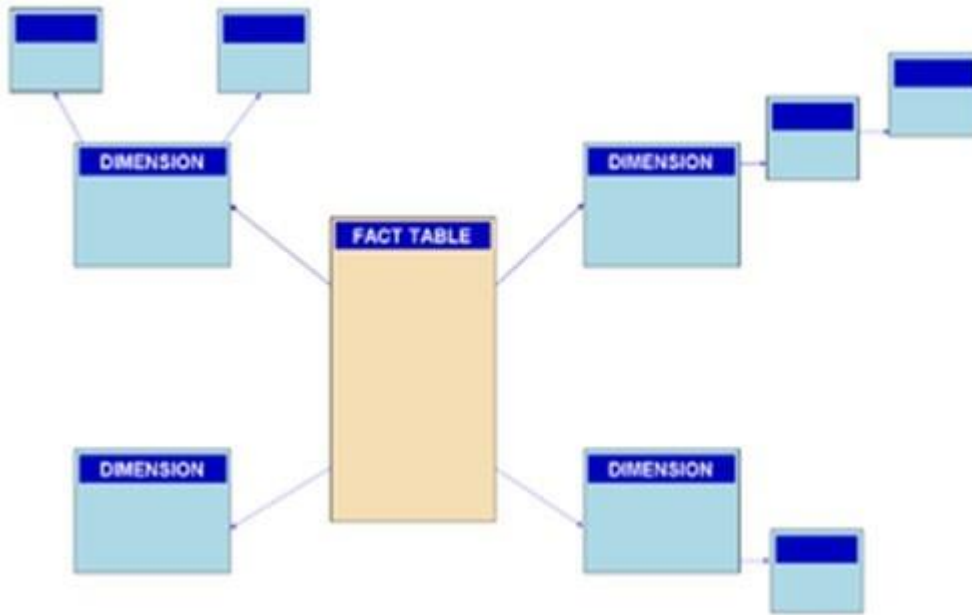


Incorrect Answers:

The data in the question is not normalized.

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

Example:



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

Box 2: dimension

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

Reference:

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

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

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

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

QUESTION 9

HOTSPOT

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

Hot Area:

Answer Area

The massively parallel processing (MPP) engine of Azure

Synapse Analytics

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

Correct Answer:

Answer Area

The massively parallel processing (MPP) engine of Azure

Synapse Analytics

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

Section: Describe core data concepts

Explanation

Explanation/Reference:

Reference:

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

QUESTION 10

HOTSPOT

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

Hot Area:

Answer Area

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

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

Correct Answer:

perform ACID transactions.

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

Incorrect Answers:

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

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

Reference:

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

QUESTION 12

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: Describe core data concepts

Explanation

Explanation/Reference:

Explanation:

- C. locally-redundant storage (LRS)
- D. geo-redundant storage (GRS)

Correct Answer: BD

Section: Describe core data concepts

Explanation

Explanation/Reference:

Explanation:

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

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

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

Reference:

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

QUESTION 17

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Explanation

Explanation/Reference:

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-sql-statements>

QUESTION 19

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

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

What should you create?

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

Correct Answer: B

Section: Describe core data concepts

Explanation

Explanation/Reference:

Explanation:

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

Reference:

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

QUESTION 20

HOTSPOT

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

Hot Area:

Answer Area

A key/value data store is optimized for

	▼
enforcing constraints	
simple lookups	
table joins	
transactions	

Correct Answer:

Answer Area

A key/value data store is optimized for

- enforcing constraints
- simple lookups
- table joins
- transactions

Section: Describe core data concepts
Explanation

Explanation/Reference:
 Explanation:

Box 1: simple lookups

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

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

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

QUESTION 21
DRAG DROP

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

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

NOTE: Each correct match is worth one point.

Select and Place:

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

Correct Answer:

Data Types	Answer Area
Image files	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

Section: Describe core data concepts
Explanation

Explanation/Reference:
Explanation:

Box 1: Image files
Azure Blob storage is suitable for image files.

Box 2: Key/value pairs
Azure CosmosDB table API is a key-value storage hosted in the cloud.

Box 3: Relationship between employees
One-to-many relationships between business domain objects occur frequently: for example, one department has many employees. There are several ways to implement one-to-many relationships in the Azure Table service.

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

QUESTION 22
DRAG DROP

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

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

NOTE: Each correct match is worth one point.

Select and Place:

Terms	Answer Area
Azure Storage account	Azure Resource Group
File share	Term
Container	Term
	Folders
	Files

Correct Answer:

Terms

Azure Storage account

File share

Container

Answer Area

Azure Resource Group

Azure Storage account

File share

Folders

Files

Section: Describe core data concepts

Explanation

Explanation/Reference:

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 23

What are two characteristics of real-time data processing? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

Correct Answer: BD

Section: Describe core data concepts

Explanation

Explanation/Reference:

Explanation:

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

Reference:

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

QUESTION 24

DRAG DROP

Match the Azure Data Factory components to the appropriate descriptions.

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

NOTE: Each correct match is worth one point.

Select and Place:

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

Correct Answer:

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

Section: Describe core data concepts

Explanation

Explanation/Reference:

Explanation:

Box 1: Dataset

Datasets must be created from paths in Azure datastores or public web URLs, for the data to be accessible by Azure Machine Learning.

Box 2: Linked service

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

Box 3: Pipeline

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

Reference:

<https://k21academy.com/microsoft-azure/dp-100/datastores-and-datasets-in-azure/>

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

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

QUESTION 25
DRAG DROP

Match the types of workloads to the appropriate scenarios.

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

NOTE: Each correct match is worth one point.

Select and Place:

Workload Types	Answer Area
Batch	Workload type
Streaming	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.

Correct Answer:

Workload Types	Answer Area
Batch	Batch
Streaming	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.
	Batch
	Updates to inventory data will be loaded to a data warehouse every 1,000 transactions.

Section: Describe core data concepts
Explanation

Explanation/Reference:
Explanation:

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

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

Box 3: Batch

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

QUESTION 26

HOTSPOT

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

Hot Area:

Answer Area

Transparent Data Encryption (TDE) encrypts

<input type="text"/>
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.

Correct Answer:

Answer Area

Transparent Data Encryption (TDE) encrypts

<input type="text"/>
a column to protect data at rest and in transit.
queries and their results in order to protect data in transit.
the database to protect data at rest.
the server to protect data at rest.

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

Reference:

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

QUESTION 27

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

Correct Answer: B

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

Reference:

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

QUESTION 28

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

- A. built-in high availability

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

Correct Answer: B

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

Reference:

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

QUESTION 29

HOTSPOT

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

Hot Area:

Answer Area

By default, each Azure SQL database is protected by

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

Correct Answer:

Answer Area

By default, each Azure SQL database is protected by

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

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

Explanation:

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

Reference:

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

QUESTION 30

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

Correct Answer: A

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

Reference:

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

QUESTION 31

DRAG DROP

Match the security components to the appropriate scenarios.

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

NOTE: Each correct match is worth one point.

Select and Place:

Components

- Authentication
- Firewall
- Encryption

Answer Area

- Prevent access to an Azure SQL database from another network.
- Support Azure Active Directory (Azure AD) sign-ins to an Azure SQL database.
- Ensure that sensitive data never appears as plain text in an Azure SQL database.

Correct Answer:

Components

- Authentication
- Firewall
- Encryption

Answer Area

- Firewall Prevent access to an Azure SQL database from another network.
- Authentication Support Azure Active Directory (Azure AD) sign-ins to an Azure SQL database.
- Encryption Ensure that sensitive data never appears as plain text in an Azure SQL database.

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

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 32

You have a transactional application that stores data in an Azure SQL managed instance.

When should you implement a read-only database replica?

- A. You need to generate reports without affecting the transactional workload.
- B. You need to audit the transactional application.
- C. You need to implement high availability in the event of a regional outage.
- D. You need to improve the recovery point objective (RPO).

Correct Answer: A

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

Explanation:

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

Reference:

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

QUESTION 33

HOTSPOT

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

Hot Area:

Answer Area

A relational database must be used when

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

Correct Answer:

Answer Area

A relational database must be used when

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

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

QUESTION 34

HOTSPOT

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

Hot Area:

Answer Area

Relational data uses

▼
collections
columns
keys
partitions

 to enforce relationships between different tables.

Correct Answer:

Answer Area

Relational data uses

▼
collections
columns
keys
partitions

 to enforce relationships between different tables.

Section: Describe how to work with relational data on Azure
Explanation

Explanation/Reference:

Reference:

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

QUESTION 35

HOTSPOT

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

Hot Area:

Answer Area

▼
A heap
A stored procedure
A view
An index

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

Correct Answer:

Answer Area

▼
A heap
A stored procedure
A view
An index

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

Section: Describe how to work with relational data on Azure
Explanation

Explanation/Reference:

Reference:

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

QUESTION 36

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.

Correct Answer: BCE

Section: Describe how to work with relational data on Azure
Explanation

Explanation/Reference:

Explanation:

Incorrect Answers:

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

Reference:

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

QUESTION 37

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

ProductName	Quantity
Product1	100
Product2	129
Product3	176

QUESTION 40

DRAG DROP

Match the terms to the appropriate descriptions.

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

NOTE: Each correct match is worth one point.

Select and Place:

Terms

- Index
- View
- Table

Answer Area

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

Correct Answer:

Terms

- Index
- View
- Table

Answer Area

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

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

Reference:

[https://en.wikipedia.org/wiki/Table_\(database\)#:~:text=A%20table%20is%20a%20collection,table%20format%20within%20a%20database.&text=In%20relational%20databases%2C%20and%20flat,a%20row%20and%20column%20intersect.](https://en.wikipedia.org/wiki/Table_(database)#:~:text=A%20table%20is%20a%20collection,table%20format%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/Database_index#:~:text=A%20database%20index%20is%20a,maintain%20the%20index%20data%20structure.

QUESTION 41

You have an e-commerce application that reads and writes data to an Azure SQL database.

Which type of processing does the application use?

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

Correct Answer: D

Section: Describe how to work with relational data on Azure

Explanation

Box 3: Azure Data Studio

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

Box 4: Microsoft SQL Server Data Tools (SSDT)

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

Reference:

<https://docs.microsoft.com/en-us/sql/ssms/download-sql-server-management-studio-ssms>

<https://code.visualstudio.com/docs/supporting/FAQ>

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

<https://docs.microsoft.com/en-us/sql/ssdt/download-sql-server-data-tools-ssdt>

QUESTION 45

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: Describe how to work with relational data on Azure**Explanation****Explanation/Reference:**

Explanation:

Box 1: Yes

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

Box 2: No

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

Box 3: Yes

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

Reference:

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

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

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

QUESTION 46

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

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

Correct Answer: A

Section: Describe how to work with relational data on Azure**Explanation****Explanation/Reference:**

Explanation:

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

Incorrect Answers:

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

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

Reference:

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

QUESTION 47

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

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 48

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: Describe how to work with relational data on Azure
Explanation

Explanation/Reference:

Reference:

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

QUESTION 50

HOTSPOT

You have the following SQL query.

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

What are `dbo.Products` and `ProductName`? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Dbo.Products :

	▼
A column	
A database	
A table	
An index	

ProductName :

	▼
A column	
A database	
A table	
An index	

Correct Answer:

Answer Area

Dbo.Products :

	▼
A column	
A database	
A table	
An index	

ProductName :

	▼
A column	
A database	
A table	
An index	

Section: Describe how to work with relational data on Azure
Explanation

Explanation/Reference:

QUESTION 51
HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

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 52

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

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

Correct Answer: D

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

Reference:

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

QUESTION 53

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

Explanation:

Box 1: Yes

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

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

Box 2: No

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

Box 3: Yes

You can use the Azure Data Studio to restore databases.

Reference:

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

QUESTION 54

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

Which Azure service should you use to support the application?

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

Correct Answer: C

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

Explanation:

Azure SQL Database is relational database and a managed service.

Incorrect Answers:

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

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

Reference:

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

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

Correct Answer: AD

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

Explanation:

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

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

Reference:

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

QUESTION 56

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: Describe how to work with relational data on Azure

Explanation

Explanation/Reference:

Explanation:

Box 1: No

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

Box 2: Yes

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

Box 3: No

Reference:

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

QUESTION 57

DRAG DROP

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

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

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

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

How should you complete the query? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Values

CREATE	GROUP BY
ORDER BY	SELECT

Answer Area

Value	SalesDate, AVG(SalesAmount)
FROM Sales	
Value	SalesDate
ORDER BY SalesDate	

QUESTION 59

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

Which type of Azure Storage should you use?

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

Correct Answer: C

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

Explanation

Explanation/Reference:

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 60

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

Correct Answer: C

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

Explanation

Explanation/Reference:

Reference:

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

QUESTION 61

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

Correct Answer: A

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

Explanation

Explanation/Reference:

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 62

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

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

Explanation

Explanation/Reference:

Explanation:

Correct Answer: C

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

Explanation

Explanation/Reference:

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 65

Your company is designing a data store for internet-connected temperature sensors.

The collected data will be used to analyze temperature trends.

Which type of data store should you use?

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

Correct Answer: B

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

Explanation

Explanation/Reference:

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 66

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Correct Answer:

Answer Area

Statements	Yes	No
When ingesting data from Azure Data Lake Storage across Azure regions, you will incur costs for bandwidth.	<input checked="" type="radio"/>	<input type="radio"/>
You can use blob, table, and file storage in the same Azure Storage account.	<input checked="" type="radio"/>	<input type="radio"/>
You implement Azure Data Lake Storage by creating an Azure Storage account.	<input checked="" type="radio"/>	<input type="radio"/>

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

Explanation/Reference:

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 67

HOTSPOT

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

Hot Area:

Answer Area

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

▼
graph.
table.
partition key.
document.

Correct Answer:

Correct Answer:

Data Store Types

- Graph
- Object
- Key/value

Answer Area

- Key/value Application users and their default language
- Object Medical images and their associated metadata
- Graph Employee data that shows the relationships between employees

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

Explanation

Explanation/Reference:

Reference:

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

QUESTION 71

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

Correct Answer: AC

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

Explanation

Explanation/Reference:

Reference:

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

QUESTION 72

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

Correct Answer: D

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

Explanation

Explanation/Reference:

QUESTION 73

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

Correct Answer: B

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

Explanation

Explanation/Reference:

Reference:

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

APIs

- Cassandra API
- Gremlin API
- MongoDB API
- Table API

Answer Area

-
-
- Key/value data

Correct Answer:

APIs

- Cassandra API
- Gremlin API
- MongoDB API
- Table API

Answer Area

- Gremlin API
- MongoDB API
- Table API

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

Explanation/Reference:

Reference:

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

QUESTION 77
HOTSPOT

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

Hot Area:

Answer Area

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

<div style="border: 1px solid black; height: 15px; width: 100%;"></div>
enable the hierarchical namespace.
set Account kind to BlobStorage.
set Performance to Premium.
set Replication to Read-access geo-redundant storage (RA-GRS).

Correct Answer:

Answer Area

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

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

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

Explanation

Explanation/Reference:

Reference:

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

QUESTION 78

HOTSPOT

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

Hot Area:

Answer Area

You can query a graph database in Azure Cosmos DB

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

Correct Answer:

Answer Area

You can query a graph database in Azure Cosmos DB

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

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

Explanation

Explanation/Reference:

Reference:

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

QUESTION 79

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

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

Correct Answer: B

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

Explanation

Explanation/Reference:

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 80

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

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

Correct Answer: D

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

Explanation

Explanation/Reference:

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 81

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

Which type of data store should you use?

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

Correct Answer: A

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

Explanation

Explanation/Reference:

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 82
HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

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

Explanation/Reference:

Explanation:

Box 1: Yes

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

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

Box 2: Yes

You can stream data into Azure Databricks using Event Hubs.

Box 3: Yes

You can run Spark jobs with data stored in Azure Cosmos DB using the Cosmos DB Spark connector. Cosmos can be used for batch and stream processing, and as a serving layer for low latency access.

You can use the connector with Azure Databricks or Azure HDInsight, which provide managed Spark clusters on Azure.

Reference:

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

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

QUESTION 83

DRAG DROP

Match the datastore services to the appropriate descriptions.

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

NOTE: Each correct match is worth one point.

Select and Place:

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

Correct Answer:

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

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

Explanation**Explanation/Reference:**

Explanation:

Box 1: Azure Cosmos DB

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

Box 2: Azure Files

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

Reference:

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

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

QUESTION 84

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

What should you use to query the files?

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

Correct Answer: C

Section: Describe an analytics workload on Azure

Explanation

Explanation/Reference:

Reference:

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

QUESTION 85

DRAG DROP

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

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

NOTE: Each correct match is worth one point.

Select and Place:

Workload Types

Batch

Streaming

Answer Area

Workload type

Analyzing web traffic from the past year

Workload type

Classifying images that were uploaded last month

Workload type

Tracking how many people are currently using a website

**Correct Answer:
Workload Types**

Batch

Streaming

Answer Area

Batch

Analyzing web traffic from the past year

Batch

Classifying images that were uploaded last month

Streaming

Tracking how many people are currently using a website

Section: Describe an analytics workload on Azure

Explanation

Explanation/Reference:

Explanation:

Box 1: Batch

The batch processing model requires a set of data that is collected over time while the stream processing model requires data to be fed into an analytics tool, often in micro-batches, and in real-time.

The batch Processing model handles a large batch of data while the Stream processing model handles individual records or micro-batches of few records.

In Batch Processing, it processes over all or most of the data but in Stream Processing, it processes over data on a rolling window or most recent record.

Box 2: Batch

Box 3: Streaming

Reference:

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

QUESTION 87

HOTSPOT

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

Hot Area:

Answer Area

In a data warehousing workload, data

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

Correct Answer:

Answer Area

In a data warehousing workload, data

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

Section: Describe an analytics workload on Azure

Explanation

Explanation/Reference:

Explanation:

Note: The data warehouse workload encompasses:

- The entire process of loading data into the warehouse
- Performing data warehouse analysis and reporting
- Managing data in the data warehouse
- Exporting data from the data warehouse

Reference:

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

QUESTION 88

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Services

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

Answer Area

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

Correct Answer:

Services

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

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

Section: Describe an analytics workload on Azure
Explanation

Explanation/Reference:
 Explanation:

Box 1: Azure Data Factory

Box 2: Azure Data Lake Storage

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

Box 3: Azure Synapse Analytics

Use Azure Synapse Analytics Workspaces.

Reference:

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

QUESTION 90
HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: Describe an analytics workload on Azure

Explanation

Explanation/Reference:

Explanation:

Box 1: Yes

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

Box 2: Yes

You can use the Azure portal to pause and resume the dedicated SQL pool compute resources. Pausing the data warehouse pauses compute. If your data warehouse was paused for the entire hour, you will not be charged compute during that hour.

Box 3: No

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

Reference:

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

QUESTION 91

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

- A. Charticator
- B. Power BI Desktop
- C. the Power BI service

D. Power BI Report Builder

Correct Answer: D

Section: Describe an analytics workload on Azure

Explanation

Explanation/Reference:

Explanation:

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

Reference:

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

QUESTION 92

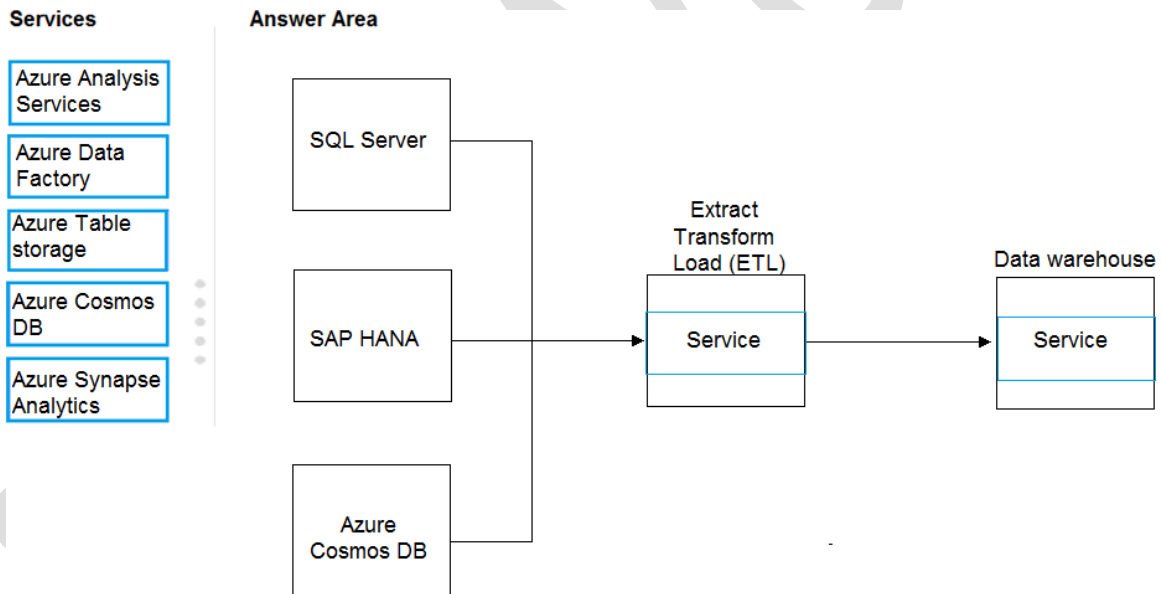
DRAG DROP

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

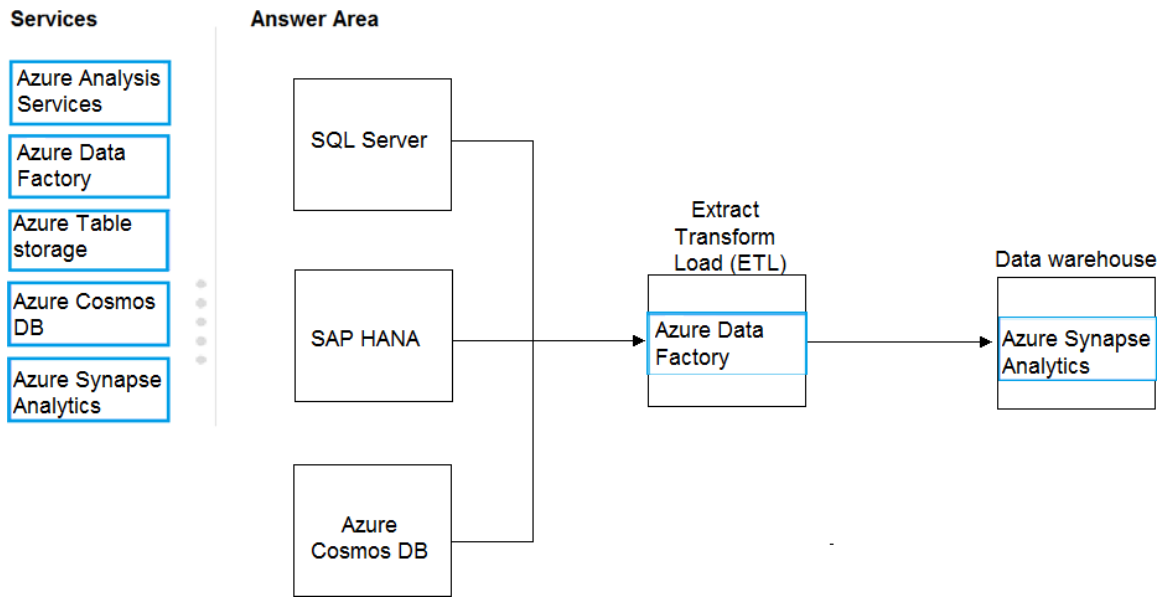
To answer, drag the appropriate service from the column on the left to its location on the right. Each service may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:



Correct Answer:



Section: Describe an analytics workload on Azure

Explanation

Explanation/Reference:

Explanation:

Box 1: Azure Data factory

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

Box 2: Azure Synapse Analytics

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

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

- Copy activity with supported source/sink matrix table
- Mapping data flow
- Lookup activity
- GetMetadata activity

Reference:

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

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

QUESTION 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.

Hot Area:

Correct Answer:

Answer Area

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

Section: Describe an analytics workload on Azure
Explanation

Explanation/Reference:

QUESTION 97

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

Correct Answer: D

Section: Describe an analytics workload on Azure
Explanation

Explanation/Reference:

Reference:

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

QUESTION 98

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

Correct Answer: B

Section: Describe an analytics workload on Azure**Explanation****Explanation/Reference:**

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-manage-compute-overview>**QUESTION 99**

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

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

Correct Answer: B**Section: Describe an analytics workload on Azure****Explanation****Explanation/Reference:**

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/concepts-pipeline-execution-triggers#trigger-execution>**QUESTION 100**

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

Which type of processing does the reporting solution use?

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

Correct Answer: C**Section: Describe an analytics workload on Azure****Explanation****Explanation/Reference:**

Reference:

<https://datawarehouseinfo.com/how-does-oltp-differ-from-olap-database/>**QUESTION 101****DRAG DROP**

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

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

NOTE: Each correct match is worth one point.**Select and Place:**

Activity Types

- Control
- Data movement
- Data transformation

Answer Area

- Copy
- Mapping data flow
- Until

Correct Answer:

Activity Types

- Control
- Data movement
- Data transformation

Answer Area

- Copy
- Mapping data flow
- Until

Section: Describe an analytics workload on Azure
Explanation

Explanation/Reference:

Reference:

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

QUESTION 102

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

Correct Answer: BDF

Section: Describe an analytics workload on Azure

Explanation

Explanation/Reference:

Explanation:

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

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

F: Transactional data tends to be highly normalized.

Reference:

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

QUESTION 103

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

Correct Answer: AD

Section: Describe an analytics workload on Azure

Explanation

Explanation/Reference:

QUESTION 104

HOTSPOT

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

Hot Area:

Answer Area

In Azure Data Factory, you can use to orchestrate pipeline

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

activities that depend on the output of other pipeline activities.

Correct Answer:

Answer Area

In Azure Data Factory, you can use to orchestrate pipeline

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

activities that depend on the output of other pipeline activities.

Section: Describe an analytics workload on Azure

Explanation

Explanation/Reference:

Reference:

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

QUESTION 105

You need to develop a solution to provide data to executives. The solution must provide an interactive graphical interface, depict various key performance indicators, and support data exploration by using drill down.

What should you use in Microsoft Power BI?

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

Correct Answer: B

Section: Describe an analytics workload on Azure

Explanation

Explanation/Reference:

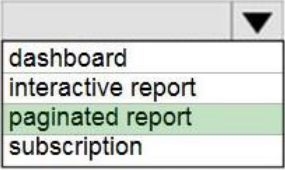
Reference:

<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>

Answer Area

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

Section: Describe an analytics workload on Azure**Explanation****Explanation/Reference:**

Explanation:

Paginated Reports in Power BI now allows users to generate these fixed-layout documents optimized for printing and archiving, such as PDF and Word files.

These document-style reports with visualizations that provide additional control, like which tables expand horizontally and vertically to display all their data and continue from page to page as needed.

Reference:

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

QUESTION 112

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

NOTE: Each correct selection is worth one point.

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

Correct Answer: AC

Section: Describe an analytics workload on Azure**Explanation****Explanation/Reference:**

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-visualizatio>

