

Microsoft.DP-900-KR.v2025-03-05.q120

□□□□:	DP-900-KR
□□□□:	Microsoft Azure Data Fundamentals (DP-900 Korean Version)
□□□:	Microsoft
□□ □□ □□□:	120
□□:	v2025-03-05
# □□ □:	526
# □□ □□□:	1200
https://www.krdump.com/Microsoft.DP-900-KR.v2025-03-05.q120.html	

NEW QUESTION: 1

□□□ □□□□□ □□□□□□ □□ □□□ □□ □□□ □□□?

- A. □□□□□□ □□□□ □□□ □□□□ □□□ □ □□□□.
- B. □□□□□□ □□□□ □□□ □□□□□ □□□ □ □□□□.
- C. □□□□□□ □□□□ □□□ □□□□□ □□ □□□□ □□□ □ □□□□.
- D. □□ □□□ □□□ □□□□□□ □□□ □□□□.

Answer: A (LEAVE A REPLY)

A pipeline itself is a group of activities that performs a task altogether in an Azure data factory. It means pipeline allows the management of activities as a set, not individually. One can not deploy and schedule activities independently, but as a pipeline only.

Option A is correct. A pipeline is a group of activities, grouped as per logic and performs a task altogether in a data factory. The pipeline allows the management of activities as a set.

Option B is incorrect. A pipeline is a group of activities and does not allow the management of activities individually.

Option C is incorrect. A pipeline allows the management of activities as a set only, not individually.

Option D is incorrect. A pipeline allows the management of activities as a set.

Reference: To know more about the Pipeline and Activities in Azure Data Factory, refer to the documentation below from Azure:

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

NEW QUESTION: 2

□□□ □□

□□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

Answer Area



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

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

Answer:
Answer Area

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

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

Explanation:
A key mechanism that allows Azure Data Lake Storage Gen2 to provide file system performance at object storage scale and prices is the addition of a hierarchical namespace. This allows the collection of objects/files within an account to be organized into a hierarchy of directories and nested subdirectories in the same way that the file system on your computer is organized. With a hierarchical namespace enabled, a storage account becomes capable of providing the scalability and cost-effectiveness of object storage, with file system semantics that are familiar to analytics engines and frameworks. One advantage of hierarchical namespace is atomic directory manipulation: Object stores approximate a directory hierarchy by adopting a convention of embedding slashes (/) in the object name to denote path segments.

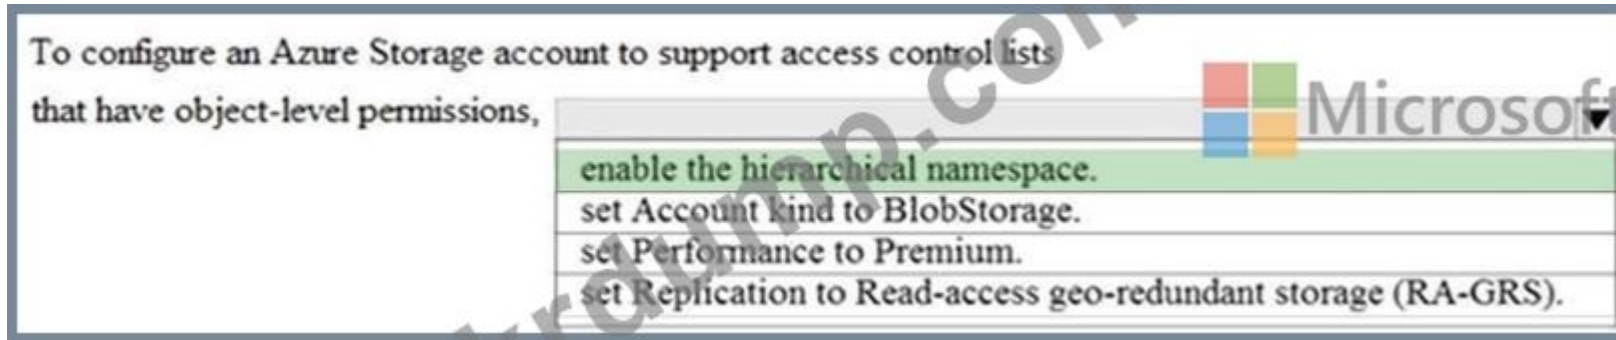
Reference:
<https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-namespace>

NEW QUESTION: 3

□□□ □□
□□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

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:

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>

NEW QUESTION: 4

Which of the following is a characteristic of a data lake?

A. Data is processed in real-time.

B. Latency in delivering data processing results is acceptable.

C. Data is always inserted one row at a time.

D. Processing can only execute serially.

Answer: D ([LEAVE A REPLY](#))

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

NEW QUESTION: 5

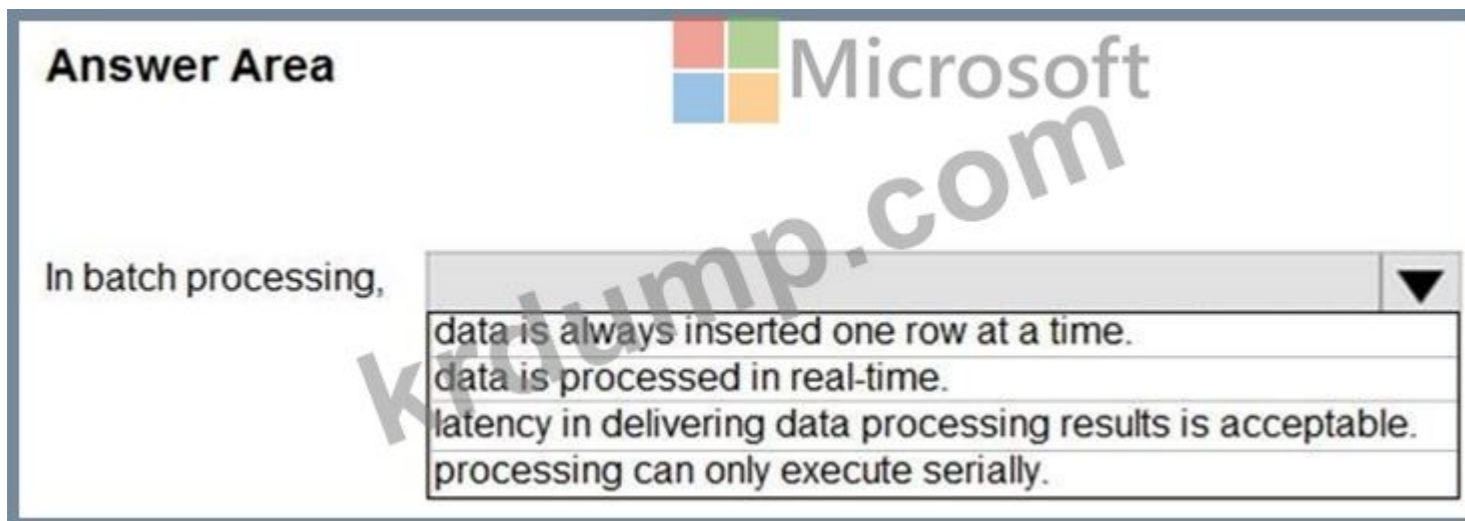
Which of the following is a characteristic of a data lake?

A. Data is processed in real-time.

B. Latency in delivering data processing results is acceptable.

C. Data is always inserted one row at a time.

D. Processing can only execute serially.



Answer:

In batch processing,

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

Explanation:

When to use batch processing.

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

Reference:

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

NEW QUESTION: 6

□□□ □□
□□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

Answer Area

Relational data is stored in

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

Answer:

Answer Area

A block of code that runs in a database is called

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

Answer:

Answer Area

A block of code that runs in a database is called

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

Explanation:

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

NEW QUESTION: 9

Which of the following is a characteristic of a data lake?

- A. It is a structured data store.
- B. It is a semi-structured data store.
- C. It is an unstructured data store.
- D. It is a hybrid data store.

Answer: (SHOW ANSWER)

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

NEW QUESTION: 10

Which of the following is a characteristic of a data lake?

It is a semi-structured data store.

- A. It is a structured data store.
- B. It is a semi-structured data store.
- C. It is an unstructured data store.
- D. It is a hybrid data store.

Answer: B,D (LEAVE A REPLY)

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>

NEW QUESTION: 11

□□□ □□
□□ □ □□□ □□ □□□□ □□ □□□□□. □□□ □□□ □□□□ □□□□□.
□□□□: □□ □□□ 1□□□□.

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

Answer:

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

Explanation:

Box 1: Yes
Use Azure Data Studio if you: Are connecting to a SQL Server 2019 big data cluster
<https://docs.microsoft.com/en-us/sql/azure-data-studio/what-is?view=sql-server-ver15> Box 2: Yes
<https://docs.microsoft.com/en-us/learn/modules/query-azure-sql-data-warehouse/4-query-dw-using-ssms> Box 3: Yes
<https://docs.microsoft.com/en-us/azure/mariadb/connect-workbench>

NEW QUESTION: 12

□□□ □□□□□□ □□ □□□ □□□□□?
A. □□ □□□ □□□□ □□□□ □□□ □□ □□ □□□□□.
B. □□ □ □□ □□□ □□□ □□ □□ □□□□ □□□□□.

C. □□□ □ □□□□ □□ □□□□ □□ □□ □□□□ □□□□□.

D. □□□□ LOB(□□ □□) □□□□□□□□ □□ □□□□□ □□□□□.

Answer: A (LEAVE A REPLY)

A data warehouse is relational in nature. This means that the structure or schema of the data is determined by predefined business and product requirements that are curated, conformed, and optimized for SQL query operations. As a result, data warehouses are best used for storing data that has been treated with a specific purpose in mind, such as data mining for BI analysis, or for sourcing a business use case that has already been identified.

<https://azure.microsoft.com/en-us/resources/cloud-computing-dictionary/what-is-a-data-warehouse/>

NEW QUESTION: 13

□□□ □□

□□□ □□□□ □□□□ □□ □□□□□.

Answer Area

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



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

Answer:

Answer Area Microsoft

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

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

Explanation:

<https://learn.microsoft.com/en-us/power-bi/create-reports/desktop-upload-desktop-files>

NEW QUESTION: 14

□□□ □□□□□ □□ □□□□ □□□ □□□ □□□□ □□□□ □□□□ □□□□. □□□□ □□□ □□ □□□ □□□□.

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

□□ □□□ □□□ □□□□ □□□□ □□□?

- A. □□□
- B. □/□
- C. □□
- D. □□□

Answer: D ([LEAVE A REPLY](#))

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

NEW QUESTION: 15

□□□ □□
□□□ □□□□ □□□□ □□ □□□□□.


Answer Area

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

dashboard
dataset
report
tile

Answer:

Answer Area

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




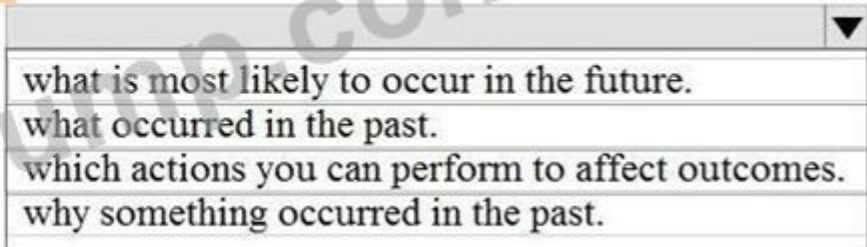
Explanation:
A Power BI report is one or more pages of visualizations, graphics, and text. All of the visualizations in a report come from a single dataset. Designers build reports and share them with others; either individually or as part of an app. Typically, Business users interact with reports in Reading view.

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


NEW QUESTION: 16

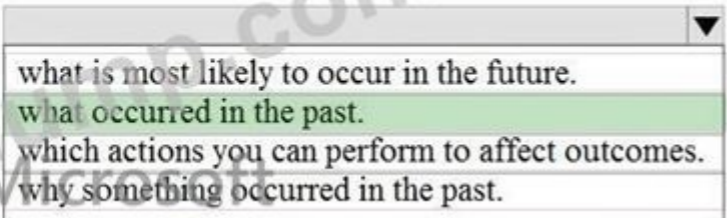
□□□ □□
□□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

Answer Area  Microsoft

Descriptive analytics tells you 

Answer:

Answer Area  Microsoft

Descriptive analytics tells you 

Explanation:
Descriptive Analytics tells you what happened in the past. Diagnostic Analytics helps you understand why something happened in the past.
Diagnostic analytics ask about the present. They drill down into why something has happened and helps users diagnose issues. Descriptive analytics ask about the past. They want to know what has been happening to the business and how this is likely to affect future sales.

- * □□□□□ □□□ □□□□ □□ □□□ □□
- * □□ □□ □□ □ □□□ □□ □□ □□□□□□.

B. * □□□ □□ □□ □□

- * □□□□ □□ □□□ □□
- * □□ □□□ □□□□ □□□ □□□ □□□□□□.

C. * □□□ □□ □□ □□

- * □□□□□ □□□□ □□□ □□□ □□□□ □□□ □□□□.
- * □□ □□□ □□□□ □□□ □□□ □□□□□□.

D. * □□□ □□ □□ □□

- * □□□□□ □□□ □□□□ □□ □□□ □□
- * □□ □□ □□ □ □□□ □□ □□ □□□ □□□□□□.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 19

- □□
- □□□□ □□□□ □□ □□□□□□.

Answer Area

The

▼
CSV
JSON
TSV
XML

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

Answer:

Answer Area

The

▼
CSV
JSON
TSV
XML

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

NEW QUESTION: 20

- Microsoft SQL Server □□□□□□□□ □□□□.
- □□□□□ □□□□□□□□□□ □□□. □□□□ □□ □□ □□□ □□□□ □□□.
- □□□□ □□□ □□□□□□□.
- □□□□□□□ □□□ □□□ □□□ □□□ □□□□□□.
- □□□ □□□□ □□□□?

A. Azure SQL Database □□ □□□□□□□

B. Azure SQL Database

C. Azure SQL

D. Azure SQL Server

Answer: C (LEAVE A REPLY)

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

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

NEW QUESTION: 21

.

Answer Area

In a SQL statement, are examples of aggregate functions.

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



Answer:

Answer Area

In a SQL statement, are examples of aggregate functions.

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



Explanation:

An aggregate function performs a calculation on a set of values, and returns a single value. The following are aggregate functions:

- * APPROX_COUNT_DISTINCT
- * AVG
- * CHECKSUM_AGG
- * COUNT
- * COUNT_BIG
- * GROUPING
- * GROUPING_ID

- * MAX
- * MIN
- * STDEV
- * STDEVP
- * STRING_AGG
- * SUM
- * VAR
- * VARP





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

NEW QUESTION: 22

□□□ □ □□ □□
 □ □□□ □□□ □□□ □□ □□□□□.
 □□□□□ □□ □□□ □□ □□□ □□□□ □□□□ □□□ □□□□.
 □ □□□ □ □ □□□ □□ □□, □□ □ □□□ □□ □□, □□ □□□□ □□ □□ □□□□.
 □□□□: □□□ □□ □□□ 1□□□□.

Tools	Answer Area		
Azure Data Studio	<table border="1"> <tr> <td data-bbox="979 1008 1072 1048">Tool</td> <td data-bbox="1289 997 2284 1068">A graphical tool for managing SQL Server or Azure SQL databases that supports access, configuration, management, and administration tasks.</td> </tr> </table>	Tool	A graphical tool for managing SQL Server or Azure SQL databases that supports access, configuration, management, and administration tasks.
Tool	A graphical tool for managing SQL Server or Azure SQL databases that supports access, configuration, management, and administration tasks.		
Microsoft SQL Server Data Tools (SSDT)	<table border="1"> <tr> <td data-bbox="979 1118 1072 1159">Tool</td> <td data-bbox="1289 1108 2284 1179">A lightweight source code editor with an mssql extension that supports connections to SQL Server and a rich editing experience for T-SQL.</td> </tr> </table>	Tool	A lightweight source code editor with an mssql extension that supports connections to SQL Server and a rich editing experience for T-SQL.
Tool	A lightweight source code editor with an mssql extension that supports connections to SQL Server and a rich editing experience for T-SQL.		
Microsoft SQL Server Management Studio (SSMS)	<table border="1"> <tr> <td data-bbox="979 1229 1072 1270">Tool</td> <td data-bbox="1289 1219 2284 1290">A lightweight editor that can run on-demand SQL queries and view and save results as text, JSON, or Microsoft Excel files.</td> </tr> </table>	Tool	A lightweight editor that can run on-demand SQL queries and view and save results as text, JSON, or Microsoft Excel files.
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	<table border="1"> <tr> <td data-bbox="979 1340 1072 1380">Tool</td> <td data-bbox="1289 1330 2284 1481">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.</td> </tr> </table>	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.
Tool	A development tool for building Azure SQL databases, Microsoft SQL Server relational databases, SQL Server Analysis Services (SSAS) data models, SQL Server Integration Services (SSIS) packages, and SQL Server Reporting Services (SSRS) reports.		

Answer:

Tools	Answer Area
	<p>Microsoft SQL Server Management Studio (SSMS) A graphical tool for managing SQL Server or Azure SQL databases that supports access, configuration, management, and administration tasks.</p>
	<p>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.</p>
	<p>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.</p>
	<p>Microsoft SQL Server Data Tools (SSDT) A development tool for building Azure SQL databases, Microsoft SQL Server relational databases, SQL Server Analysis Services (SSAS) data models, SQL Server Integration Services (SSIS) packages, and SQL Server Reporting Services (SSRS) reports.</p>

Explanation:

Box 1: Microsoft SQL Server Management Studio (SSMS)

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

Box 2: Microsoft Visual Studio Code

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>

NEW QUESTION: 23

Which of the following is a valid T-SQL query to retrieve the names of all tables in the AdventureWorks database? (Select all that apply.)

- A. `SELECT * FROM sys.tables`
- B. `SELECT * FROM sys.databases`
- C. `SELECT * FROM sys.tables WHERE database_id = 10`
- D. `SELECT * FROM sys.tables WHERE database_id = 10 AND schema_id = 1`

Answer: (SHOW ANSWER)

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>

NEW QUESTION: 24

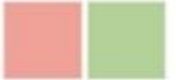

□□□ □□

□□ □ □□□ □□ □□□ □□□□□ □□ □□□□□. □□□ □□□ □□□□ □□□□□.

□□□□: □□ □□□ 1□□□□.

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

Answer:

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

Explanation:

Box 1: Yes

Batch processing refers to processing of high volume of data in batch within a specific time span.

It processes large volume of data all at once. Batch processing is used when data size is known and finite. It takes little longer time to processes data.

Box 2: Yes

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

Box 3: No

Reference:

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

NEW QUESTION: 25

Microsoft SQL Server 2019 Enterprise Edition is not supported on Azure SQL Database. Which Azure SQL service should you use to run Microsoft SQL Server 2019 Enterprise Edition?

- A. Azure SQL Database
- B. MySQL on Azure
- C. Azure SQL Managed Instance
- D. Azure SQL Database Flexible Server

Answer: C (LEAVE A REPLY)

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

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

Reference:

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

NEW QUESTION: 26

Which visualization type is best for displaying the relationship between two numerical values?

Scatter plot. Treemap. Key influencer. Visualization. A chart of colored, nested rectangles that displays individual data points represented by the size and color of a relative rectangle. A chart that displays the major contributors of a selected result or value. A chart that shows the relationship between two numerical values.


Visualizations	Answer Area
Key influencer	Visualization
Scatter	Visualization
Treemap	Visualization

A chart of colored, nested rectangles that displays individual data points represented by the size and color of a relative rectangle.
 A chart that displays the major contributors of a selected result or value.
 A chart that shows the relationship between two numerical values.

Answer:

Visualizations

Answer Area



Treemap	A chart of colored, nested rectangles that displays individual data points represented by the size and color of a relative rectangle.
Key influencer	A chart that displays the major contributors of a selected result or value.
Scatter	A chart that shows the relationship between two numerical values.

Explanation:

Box 1: Tree map

Treemaps are chart 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.


<https://docs.microsoft.com/en-us/learn/modules/explore-concepts-of-data-analytics/3-explore-data-visualization>

NEW QUESTION: 27

□□ □□□ □□ □□□ □□ □□ □□□ □□□□ □□□□ □□□.

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

Core (SQL)
Gremlin
MongoDB
Table



Explanation:

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

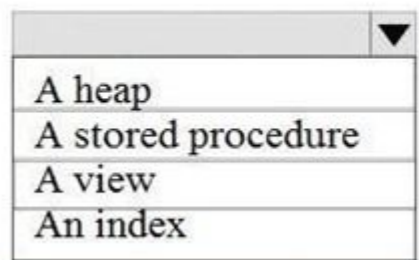
Reference:

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


NEW QUESTION: 29

□□□ □□
□□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

Answer Area

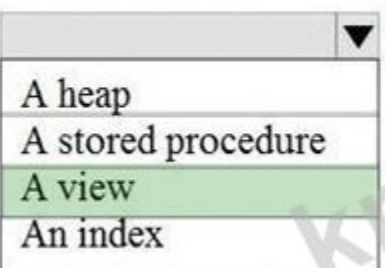
 is a virtual table that contains content defined by a query.

A heap
A stored procedure
A view
An index




Answer:

Answer Area

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

A heap
A stored procedure
A view
An index



Explanation:

Heap: tables: without a clustered index.

Stored procedure: a sub-routine available to access application in RDBMS A view: is a table that is being created based on a user query An index: a data structure that improves the speed of data retrieval operations

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

NEW QUESTION: 30

□□□ □□

□□ □ □□□ □□ □□□□ □□ □□□□□. □□□ □□□ □□□□ □□□□□.

□□□□: □□ □□□ 1□□□□.

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"/>

Answer:

Answer Area



Microsoft
Statements

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"/>

Explanation:

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

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

Azure SQL Managed Instance.

SQL Server in Azure Virtual Machines.

SQL Server Integration Services.

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

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

NEW QUESTION: 31

□□□ □□

□□□ □□□□ □□□□ □□ □□□□□.

Answer Area

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



Dropdown menu options:

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

Answer:

Answer Area

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



Dropdown menu options:

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

DP-900-KR ☐☐ ☐☐☐ ☐☐☐☐☐ ☐☐ DumpTop ☐☐ ☐☐☐☐ ☐☐☐ DP-900-KR ☐☐! DumpTop ☐ ☐☐ DP-900-KR ☐☐ ☐☐☐ ☐☐☐☐☐☐, DumpTop DP-900-KR ☐☐ ☐☐☐ ☐☐☐☐☐☐☐☐☐☐ ☐☐☐☐☐☐☐☐☐☐. ☐☐☐☐☐ ☐☐☐☐☐ ☐☐☐☐☐ ☐☐ DumpTop DP-900-KR ☐☐☐☐☐☐☐☐. <https://www.dumptop.com/Microsoft/DP-900-KR-dump.html> (316 Q&As Dumps, **30%OFF Special Discount: KrDump**)

NEW QUESTION: 32

POSIX ☐☐ ACL(☐☐☐☐ ☐☐ ☐☐)☐ ☐☐ ☐☐☐☐☐ ☐☐☐☐☐ Azure ☐☐☐☐☐☐☐☐☐☐?

- A. Azure ☐☐☐☐☐☐
- B. Azure ☐☐☐☐☐☐☐☐☐
- C. Azure ☐☐☐☐☐☐☐☐
- D. Azure ☐☐☐☐☐☐

Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 33

☐☐☐☐☐☐☐☐☐

□□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

Answer Area



A relational database must be used when

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

Answer:

Answer Area

A relational database must be used when



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

NEW QUESTION: 34

□□□ □□□□ □□(OLTP) □□□□□ □ □□ □□□ □□□□□? □ □□□ □□□ □□□□ □□□□□.

□□□□: □□ □□□ 1□□□□.

- A. □□□□□ □□□
- B. □□□□□□□ □□□ □□□
- C. □□□□□□□□ □□□ □□□□ □□ □ □□□ □□□
- D. □□□□ □□□
- E. □□□ □□□ □□□ □□
- F. □□ □□ □ □□ □□

Answer: [\(SHOW ANSWER\)](#)

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

D: Transactional data tends to be highly normalized.

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

Reference:

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

NEW QUESTION: 35

□□□ □ □□ □□

Azure SQL □□□□ □□□ □□ □□□ □□□□□. □□□□□ □□ □□□ □□□ □□□□ □□□□ □□ □□□ □□□ □□□□ □□ □ □□□□ □□ □ □□□□.

□□□□: □□□ □□ □□□ 1□□□□.

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

Answer:

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

Explanation:

Box 1: Azure SQL Database single database

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

Scenarios well suited for serverless compute

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

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

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

Box 2: Azure SQL Managed Instance

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

Box 3: Azure SQL Database elastic pool

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

Reference:

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

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

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

NEW QUESTION: 36

Which of the following is not a data manipulation language (DML) statement?

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

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 37

Which of the following is not a data manipulation language (DML) statement?

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

Answer: D ([LEAVE A REPLY](#))

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>

NEW QUESTION: 38

Microsoft Power BI Desktop is a self-hosted application. Which of the following is not a data manipulation language (DML) statement? (200 characters) Answer: INSERT

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

Answer: ([SHOW ANSWER](#))

As data acquisition and preparation is done at the Power BI Desktop side, anything related to sharing is on Power BI service.

<https://learn.microsoft.com/en-us/power-bi/fundamentals/media/service-service-vs-desktop/power-bi-venn-desktop-service.png>

NEW QUESTION: 39

Which of the following is not a data manipulation language (DML) statement?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 40

□□□ □□
□□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

Answer Area



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

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

Answer:

Answer Area

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



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

Explanation:

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

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

Incorrect Answers:

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

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


Reference:

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

NEW QUESTION: 41

□□□ □□
□□□ □□□□ □□□□ □□ □□□□□.

Answer Area



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

- Data analysts
- Data engineers
- Data scientists
- Database administrators

Answer:

Answer Area

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

- Data analysts
- Data engineers
- Data scientists
- Database administrators

NEW QUESTION: 42

External data store helps in searching the information in other types of data stores. Here in the given scenario, the primary index of the file is file path but you want to search data on the basis of the contents of the file. So, External Index Data Store will create an external index i.e. secondary index for the data store and perform the search.

- A. Document data store
- B. Key/Value data store
- C. External Index data store
- D. Graph data store

Answer: C (LEAVE A REPLY)

External data store helps in searching the information in other types of data stores. Here in the given scenario, the primary index of the file is file path but you want to search data on the basis of the contents of the file. So, External Index Data Store will create an external index i.e. secondary index for the data store and perform the search.

- Option A is incorrect. Document data store does not provide the ability to search any information that is present in any other data store.
- Option B is incorrect. Key/Value data store does not provide the ability to search any information that is present in any other data store.
- Option C is correct. External Index data store does not provide the ability to search for any information that is present in any other data store.
- Option D is incorrect. Graph data store does not provide the ability to search any information that is present in any other data store.

Reference: To know more about the different data stores, refer to the link below:
<https://docs.microsoft.com/en-us/azure/architecture/data-guide/big-data/non-relational-data>

NEW QUESTION: 43

ACID stands for Atomicity, Consistency, Isolation, and Durability. In a database system, ACID properties are used to ensure the reliability of the data. Which of the following is not a property of ACID?

- A. Atomicity
- B. Consistency
- C. Isolation
- D. Durability

Answer: (SHOW ANSWER)

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

NEW QUESTION: 44

Which of the following is not a property of ACID?

- A. Atomicity

Answer Area

Statements	Yes	No
Data stored in the Hot access tier of Azure Blob storage is stored on high-performance media.	<input checked="" type="radio"/>	<input type="radio"/>
The Cool access tier in Azure Blob storage incurs higher storage costs compared to the Hot access tier.	<input type="radio"/>	<input checked="" type="radio"/>
The Archive access tier in Azure Blob storage has the highest latency of the three available tiers.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:
<https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview> The archive tier is an offline tier for storing data that is rarely accessed. The archive access tier has the lowest storage cost. However, this tier has higher data retrieval costs with a higher latency as compared to the hot and cool tiers.

NEW QUESTION: 46

□□□ □□
 □□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

Answer Area

Relational data uses to enforce relationships between different tables.

collections
columns
keys
partitions

Answer:
Answer Area

Relational data uses to enforce relationships between different tables.

collections
columns
keys
partitions

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

DP-900-KR ☐☐ ☐☐☐ ☐☐☐☐☐ ☐☐ DumpTop ☐☐ ☐☐☐☐ ☐☐☐ DP-900-KR ☐☐! DumpTop ☐ ☐☐ DP-900-KR ☐☐ ☐☐☐ ☐☐☐☐☐☐, DumpTop DP-900-KR ☐☐ ☐☐☐ ☐☐☐☐☐☐☐☐☐ ☐☐☐ ☐☐☐☐☐☐. ☐☐☐☐☐ ☐☐☐☐☐ ☐☐☐☐☐ ☐☐ DumpTop DP-900-KR ☐☐☐ ☐☐☐☐☐. <https://www.dumptop.com/Microsoft/DP-900-KR-dump.html> (316 Q&As Dumps, 30%OFF Special Discount: KrDump)

NEW QUESTION: 47

☐☐☐ ☐☐
☐☐☐ ☐☐☐☐☐ ☐☐☐☐☐ ☐☐ ☐☐☐☐☐.

Answer Area

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

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

Answer:
Answer Area

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

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

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

NEW QUESTION: 48

☐☐☐ ☐☐
☐☐☐ ☐☐☐☐☐☐ ☐☐☐☐☐☐ ☐☐☐ ☐☐☐☐☐☐☐.

Answer Area

By default, each Azure SQL database is protected by

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

Answer:

Answer Area

By default, each Azure SQL database is protected by



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

Explanation:

When you create a new server in Azure SQL Database or Azure Synapse Analytics named mysqlserver, for example, a server-level firewall blocks all access to the public endpoint for the server Reference: <https://docs.microsoft.com/en-us/azure/security/fundamentals/infrastructure-sql>

NEW QUESTION: 49

□□□ □□□ □□ 3□□ □□□ □□□□□□.

□□ □□□□ □□□ □□ □□□□.

Sales by Quarter



4□□□□ □□ □□□ □□□ □□□□□□?

- A. □□
- B. □□□
- C. □□□
- D. □□

Answer: (SHOW ANSWER)

Predictive, to answer the question: What will happen?



Reference:

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

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

NEW QUESTION: 50

□□□ □□
 □□□ □□□□ □□□□ □□ □□□□□.

Answer Area

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

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

Answer:

Answer Area

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

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

NEW QUESTION: 51

□□□ □□
 □□ □ □□□ □□ □□□ □□□□□ □□ □□□□□. □□□ □□□ □□□□ □□□□□.

□□□□: □□ □□□ 1□□□□.

Statements

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

Answer:

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

Explanation:

Box 1: No

You can duplicate a dashboard. The duplicate ends up in the same Power BI workspace. There is no current functionality that allows you to move reports from one workspace to another.

Box 2: No

Box 3: Yes

Reference:

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

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

<https://powerbi.microsoft.com/en-us/excel-and-power-bi/>

NEW QUESTION: 52

Azure Storage □□□ Azure Data Lake Storage□ □□□□□□□□ □□□ □□□□□?

- A. □□ □□□ □□□□□ □□□.
- B. □□□ □□□□□□□□ □□□□□ □□□.
- C. □□□ □□□□□□□□ □□□□□□ □□□.
- D. □□ □□□ □□□□□□ □□□.

Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 53

□□□ □□
□□□ □□□□ □□□□ □□ □□□□□.

Answer Area



Relational databases are commonly used to store and query

blobs.
semi-structured data.
structured data.
unstructured data.

Answer:

Answer Area



Relational databases are commonly used to store and query

blobs.
semi-structured data.
structured data.
unstructured data.

NEW QUESTION: 54

Which of the following is a data warehouse?

Which of the following is a data warehouse?

- A. OLAP
- B. ETL
- C. OLTP
- D. Data Lake

Answer: C (LEAVE A REPLY)

NEW QUESTION: 55

Which of the following is a data warehouse?

Which of the following is a data warehouse?

Answer Area



Backing up and restoring data are the responsibilities of a

data analyst
data engineer
data scientist
database administrator

Answer:

- SSIS package execution

Reference:

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

NEW QUESTION: 58

□□□ □ □□ □□

□□□ □□□ □□□ □□□□□.

□□□□□ □□ □□□ □□□ □□□ □□□□ □□□ □□□□.

□ □□□ □ □ □□□ □□ □□, □ □ □□ □□□ □□ □□, □□ □□□□ □□ □□ □□□□.

□□□□: □□□ □□ □□□ 1□□□□.

Terms

- Index
- View
- Table

Answer Area

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

Answer:

Terms	Answer Area
	<input type="text"/>
	<input type="text"/>
	<input type="text"/>

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

Explanation:

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

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

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

NEW QUESTION: 59

Azure Cosmos DB □□□ □□□□□ □ Azure □□ □□□ □□□□ □□□□ □□□ □□□□□?

- A. □□ □□□ □□
- B. □□□ □□
- C. □□ □□ □□
- D. □□□ □□□ □□

Answer: B (LEAVE A REPLY)

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>

NEW QUESTION: 60

□□□ □ □□ □□

□□ □□□ □□□ □□□ □□ □□□□□.

□□□□□ □□ □□□ □□ □□□ □□□ □□□□ □□□ □□□□ □□ □ □□□□.

□□□□: □□□ □□ □□□ 1□□□□.

Job role

Data analyst

Data engineer

Database administrator

Answer Area

Job role

Database backup and recovery

Job role

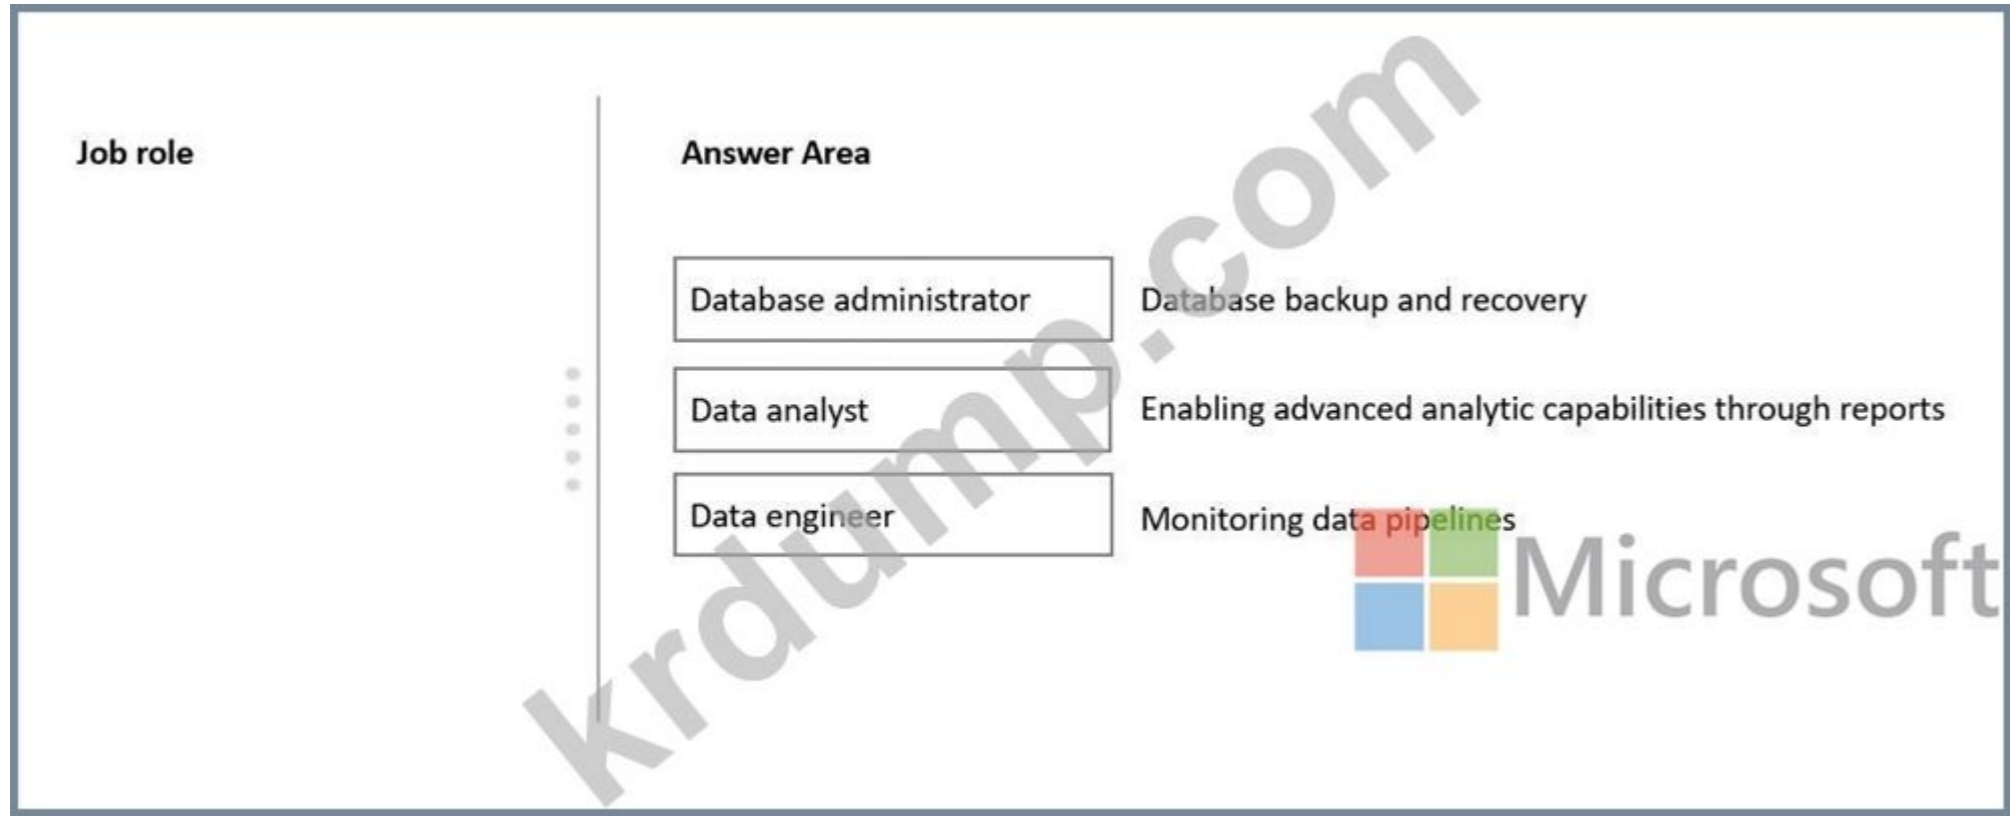
Enabling advanced analytic capabilities through reports

Job role

Monitoring data pipelines



Answer:



NEW QUESTION: 61

□□□ □□
 □□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

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

Microsoft

Answer:

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

Microsoft

Explanation:

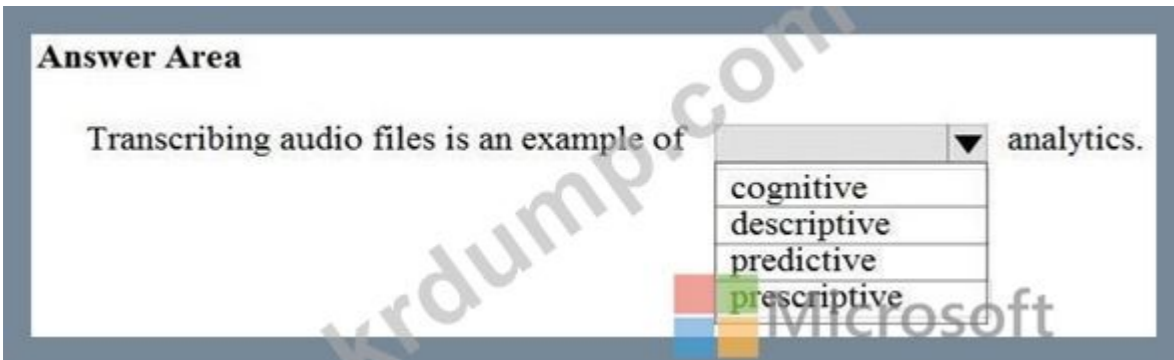
Clustered indexes sort and store the data rows in the table or view based on their key values.

These are the columns included in the index definition. There can be only one clustered index per table, because the data rows themselves can be stored in only one order.

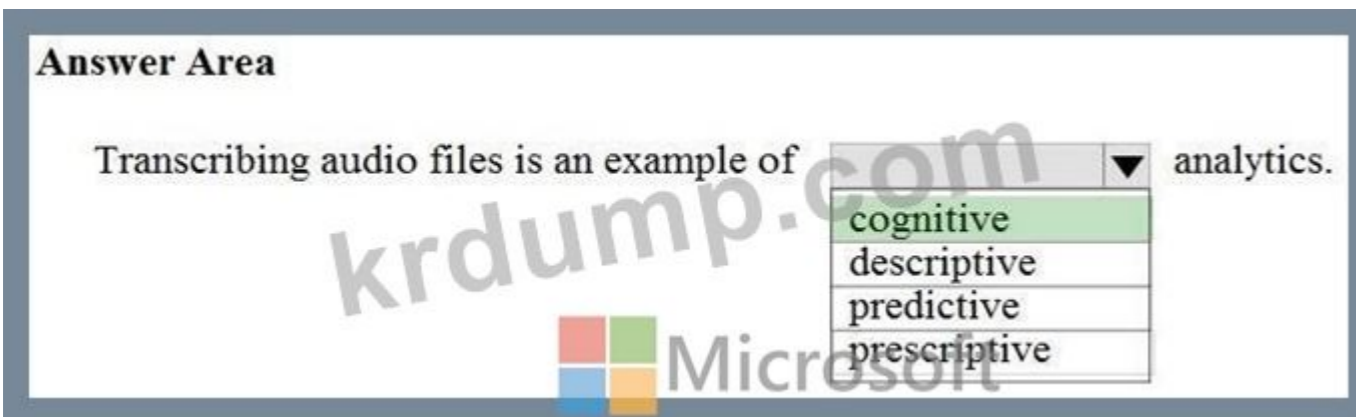
DP-900-KR [DumpTop DP-900-KR](https://www.dumptop.com/Microsoft/DP-900-KR-dump.html)! DumpTop [DP-900-KR](https://www.dumptop.com/Microsoft/DP-900-KR-dump.html) [DumpTop DP-900-KR](https://www.dumptop.com/Microsoft/DP-900-KR-dump.html) [DumpTop DP-900-KR](https://www.dumptop.com/Microsoft/DP-900-KR-dump.html) [DumpTop DP-900-KR](https://www.dumptop.com/Microsoft/DP-900-KR-dump.html) [DumpTop DP-900-KR](https://www.dumptop.com/Microsoft/DP-900-KR-dump.html). [DumpTop DP-900-KR](https://www.dumptop.com/Microsoft/DP-900-KR-dump.html) [DumpTop DP-900-KR](https://www.dumptop.com/Microsoft/DP-900-KR-dump.html). <https://www.dumptop.com/Microsoft/DP-900-KR-dump.html> (316 Q&As Dumps, **30%OFF Special Discount: KrDump**)

NEW QUESTION: 62

Transcribing audio files is an example of _____ analytics.



Answer:



Explanation:

Transcribe means converting from Audio to Text. This is similar to Alexa, Alexa is a Cognitive device.
<https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/batch-transcription>

NEW QUESTION: 63

Azure SQL _____ to offload read-only query workloads.

A. _____

B. _____

C. _____

D. _____(RPO) _____

Answer: A (LEAVE A REPLY)

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

Reference:

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

NEW QUESTION: 64

□□□ □□
□□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

Answer Area

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

- cognitive
- descriptive
- predictive
- prescriptive

Answer:

Answer Area

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

- cognitive
- descriptive
- predictive
- prescriptive

Explanation:

Speech service, part of Azure Cognitive Services, is certified by SOC, FedRAMP, PCI DSS, HIPAA, HITECH, and ISO.

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

NEW QUESTION: 65

□□□ □□
□□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

Answer Area

You can query a graph database in Azure Cosmos DB

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

Answer:

Answer Area

You can query a graph database in Azure Cosmos DB

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

Explanation:

You can use the Gremlin language to create graph entities (vertices and edges), modify properties within those entities, perform queries and traversals, and delete entities.

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

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

NEW QUESTION: 66

□□□ □ □□ □□

□□□ □□□□□ □□ □□□□□ □□□□□.

□□□□□ □□ □□□ □□ □□□□□ □□□ □□□□□ □□□ □□□□.

□ □□□□□ □ □, □ □ □□ □□ □□ □□□□ □□ □ □□□□. □□: □□□ 1□□□□.

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:

Processes



Microsoft

Answer Area

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

Extract, transform, and load (ETL) processing

Data is saved to a target system, and then processed

Extract, load, and transform (ELT) processing

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

Data ingestion

NEW QUESTION: 67

Progress indicator: 100% complete

Answer Area



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

- cognitive
- descriptive
- predictive
- prescriptive

analytics.

Answer:

Answer Area

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



	▼
cognitive	
descriptive	
predictive	
prescriptive	

Explanation:

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

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

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

Predictive, to answer the question: What will happen?

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

Reference:

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

NEW QUESTION: 68

Which of the following is a common data format?

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

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 69

Which of the following is a valid T-SQL command?

- A. SELECT
- B. INSERT
- C. UPDATE

Answer: A ([LEAVE A REPLY](#))

Table hints override the default behavior of the Query Optimizer for the duration of the data manipulation language (DML) statement by specifying a locking method, one or more indexes, a query-processing operation such as a table scan or index seek, or other options. Table hints are specified in the FROM clause of the DML statement and affect only the table or view referenced in that clause.

This applies to the following commands:

- DELETE
- INSERT
- SELECT
- UPDATE

MERGE

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/queries/hints-transact-sql-table>

NEW QUESTION: 70

Azure SQL `Products` table is queried.

Which of the following is a valid query? (3 correct answers.)

- A. `SELECT * FROM Products`
- B. `SELECT * FROM Products WHERE 1=1`
- C. `SELECT * FROM Products WITH NOLOCK`
- D. `SELECT * FROM Products WITH (NOLOCK)`
- E. `SELECT * FROM Products WITH (READ UNCOMMITTED)`

Answer: [\(SHOW ANSWER\)](#)

Incorrect Answers:

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

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

NEW QUESTION: 71

Which of the following is a valid query? (3 correct answers.)

Which of the following is a valid query? (3 correct answers.)

- A. `SELECT * FROM Products`
- B. `SELECT * FROM Products WHERE 1=1`
- C. `SELECT * FROM Products WITH NOLOCK`
- D. `SELECT * FROM Products WITH (NOLOCK)`

Answer: [B \(LEAVE A REPLY\)](#)

NEW QUESTION: 72

Which of the following is a valid query? (3 correct answers.)

Which of the following is a valid query? (3 correct answers.)



Answer:

Answer Area

Microsoft SQL Server database tables store



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

NEW QUESTION: 73

□□□ □□
 □□□ □□□□ □□□□ □□ □□□□□.

Answer Area

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



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

Answer:

Answer Area

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



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

Explanation:
 For some semi-structured data formats (e.g. JSON), data sets are frequently a simple concatenation of multiple documents.
 Reference:
<https://docs.snowflake.com/en/user-guide/semistructured-considerations.html>

NEW QUESTION: 74

□□□ □□
 □□ □ □□□ □□ □□□ □□□□□ □□ □□□□□. □□□ □□□ □□□□ □□□□□.
 □□□□: □□ □□□ 1□□□□.

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

Answer:

Answer Area



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

Explanation:

Cloud features such as scalability, high-availability and multi-tenant capability are included, reducing the amount of coding that developers must do.

Single databases in Azure SQL Database support manual dynamic scalability, but not autoscale.

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

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

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

NEW QUESTION: 75

□□□ □□

□□ □ □□□ □□ □□□ □□□□□ □□ □□□□□. □□□ □□□ □□□□ □□□□□.

□□□□: □□ □□□ 1□□□□.

☐☐ ☐☐☐ ☐☐ ☐☐☐☐☐ ☐☐ ☐☐☐☐☐. ☐☐☐ ☐☐☐ ☐☐☐☐☐☐☐☐.
☐☐☐☐: ☐☐ ☐☐☐ 1☐☐☐☐☐.

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

Answer:

Answer Area

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

Explanation:

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

NEW QUESTION: 78

☐☐☐ ☐☐☐ ☐☐☐☐☐ ☐☐☐☐☐☐☐☐ ☐☐☐ ☐☐☐☐☐☐☐☐☐☐☐. ☐☐☐☐☐☐☐☐☐☐?

- A. Azure Table Storage
- B. Azure Queue Storage
- C. Azure ☐☐

Answer: A (LEAVE A REPLY)

You can store logs in the Azure table storage or Blob storage but Microsoft itself recommends using Blob storage. Azure Storage Analytics stores log data in Blob storage. This 'Azure Storage Table Design Guide' points out several draw backs of using table storage for logs and also provides details on how to use the blob storage to store logs. Read the 'Log data anti-pattern' section in particular for this use case.

Blob storage is not an option in answers, so the choice shall be Azure table storage.
<https://learn.microsoft.com/en-us/azure/storage/tables/table-storage-design-patterns#log-data-anti-pattern>

NEW QUESTION: 79

□□□ □ □□ □□

Sales□□ □□□□ □□□, □ □□□□□ □□ □□□□ □□□□ □□□□.

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

□□ □□ □□□ □□□□□ □□□□ □□□□ □□□□ □□□□ □□□□.

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

□□□ □□□ □□□□ □□□? □□□□□ □□□ □□ □□□ □□□□□□□□. □ □□ □ □, □ □ □□ □□ □□ □□□□ □□ □ □□□□. □□ □□□ □□ □□□ □□□□□□ □□□□□ □□□□ □ □ □□ □□.

□□□□: □□ □□□ 1□□□□.

Values

CREATE

GROUP BY

ORDER BY

SELECT

Answer Area

Value

 SalesDate, AVG(SalesAmount)

FROM Sales

Value

 SalesDate

ORDER BY SalesDate

Answer:

Values


CREATE	GROUP BY
ORDER BY	SELECT

Answer Area

```

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

```



Explanation:
 Box 1: SELECT
 Box 2: GROUP BY
 Example:
 When used with a GROUP BY clause, each aggregate function produces a single value covering each group, instead of a single value covering the whole table. The following example produces summary values for each sales territory in the AdventureWorks2012 database. The summary lists the average bonus received by the sales people in each territory, and the sum of year-to-date sales for each territory.
 SELECT TerritoryID, AVG(Bonus)as 'Average bonus', SUM(SalesYTD) as 'YTD sales' FROM Sales.SalesPerson GROUP BY TerritoryID; Reference:
<https://docs.microsoft.com/en-us/sql/t-sql/functions/avg-transact-sql>

NEW QUESTION: 80

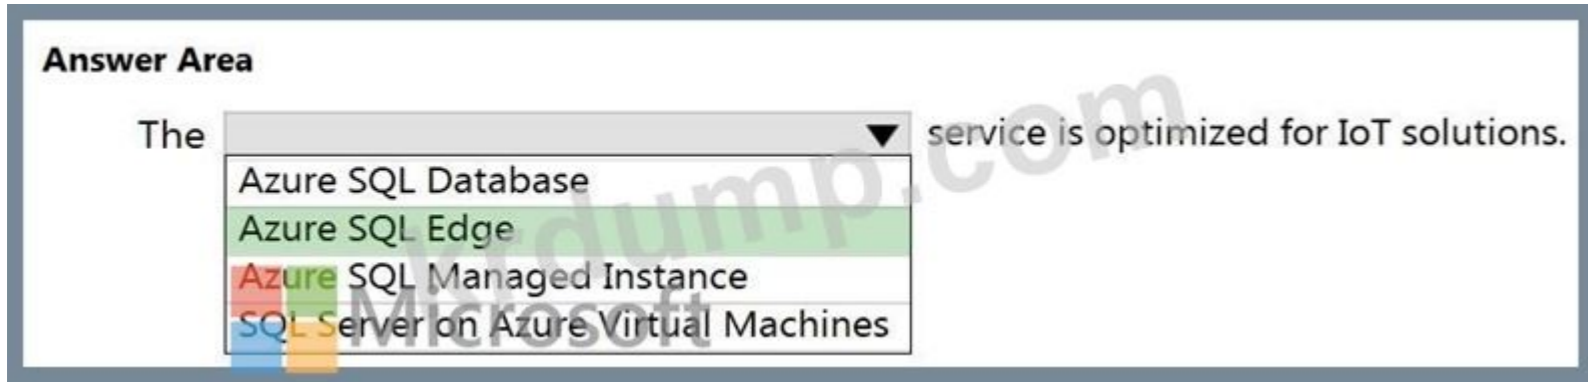
□□□ □□
 □□□ □□□□ □□□□ □□ □□□□□.

Answer Area

The service is optimized for IoT solutions.

- Azure SQL Database
- Azure SQL Edge
- Azure SQL Managed Instance
- SQL Server on Azure Virtual Machines

Answer:



NEW QUESTION: 81

- Which of the following is a cross-platform database tool for data professionals who use on-premises and cloud data platforms on Windows, macOS, and Linux?
- A. Azure Data Studio
 - B. sqlcmd
 - C. Microsoft SQL Server Enterprise Edition (SSMS)
 - D. Azure SQL Database

Answer: A (LEAVE A REPLY)

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>

NEW QUESTION: 82

- Which of the following are properties of Azure Table storage?
- A. It is a NoSQL database.
 - B. It is a key-value store.
 - C. It is a document database.
 - D. It is a columnar database.

Answer: A,C (LEAVE A REPLY)

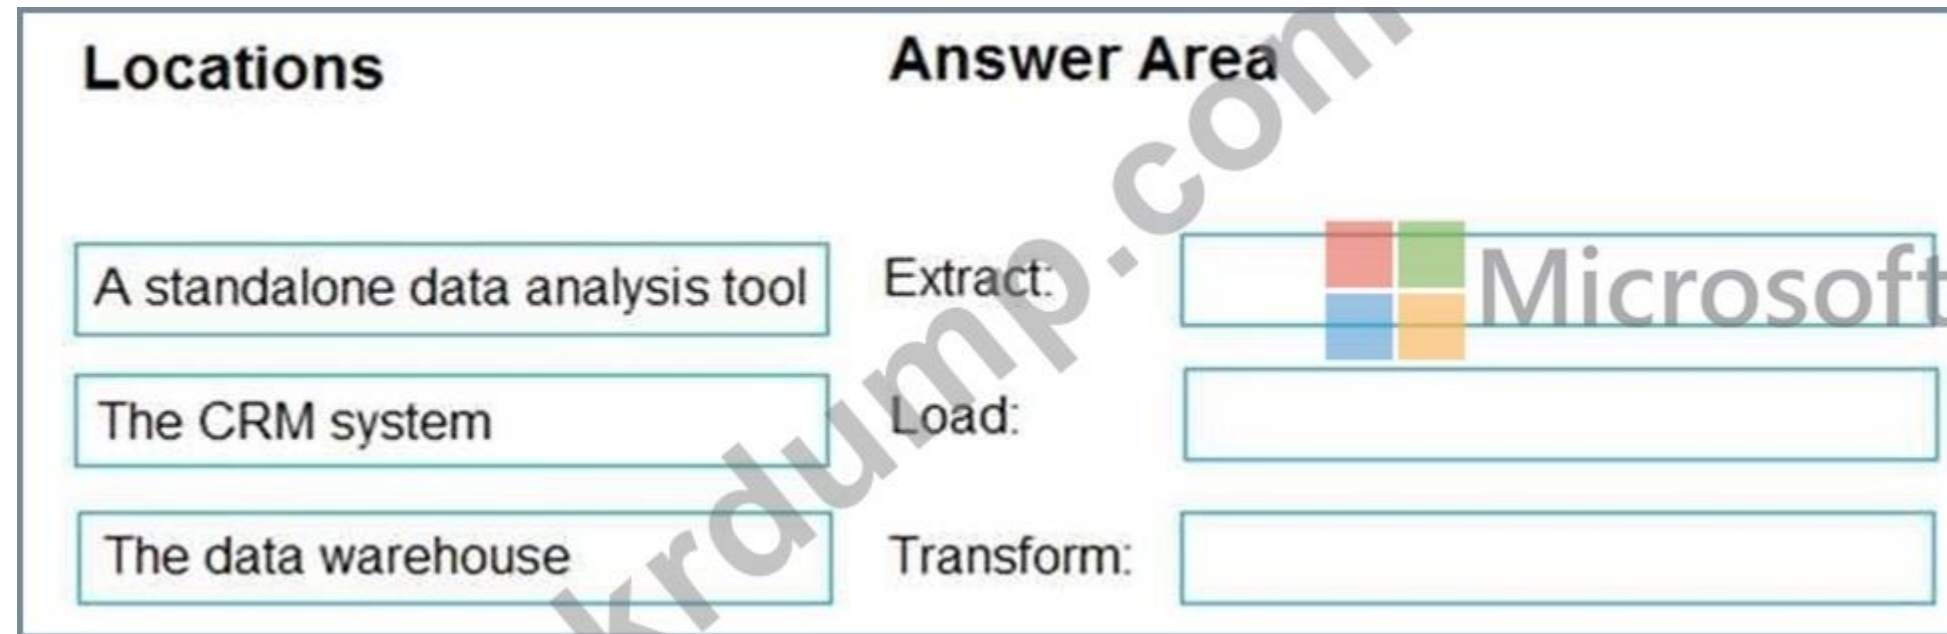
Properties: A property is a name-value pair. Each entity can include up to 252 properties to store data. Each entity also has three system properties that specify a partition key, a row key, and a timestamp. Entities with the same partition key can be queried more quickly, and inserted/updated in atomic operations. An entity's row key is its unique identifier within a partition.

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

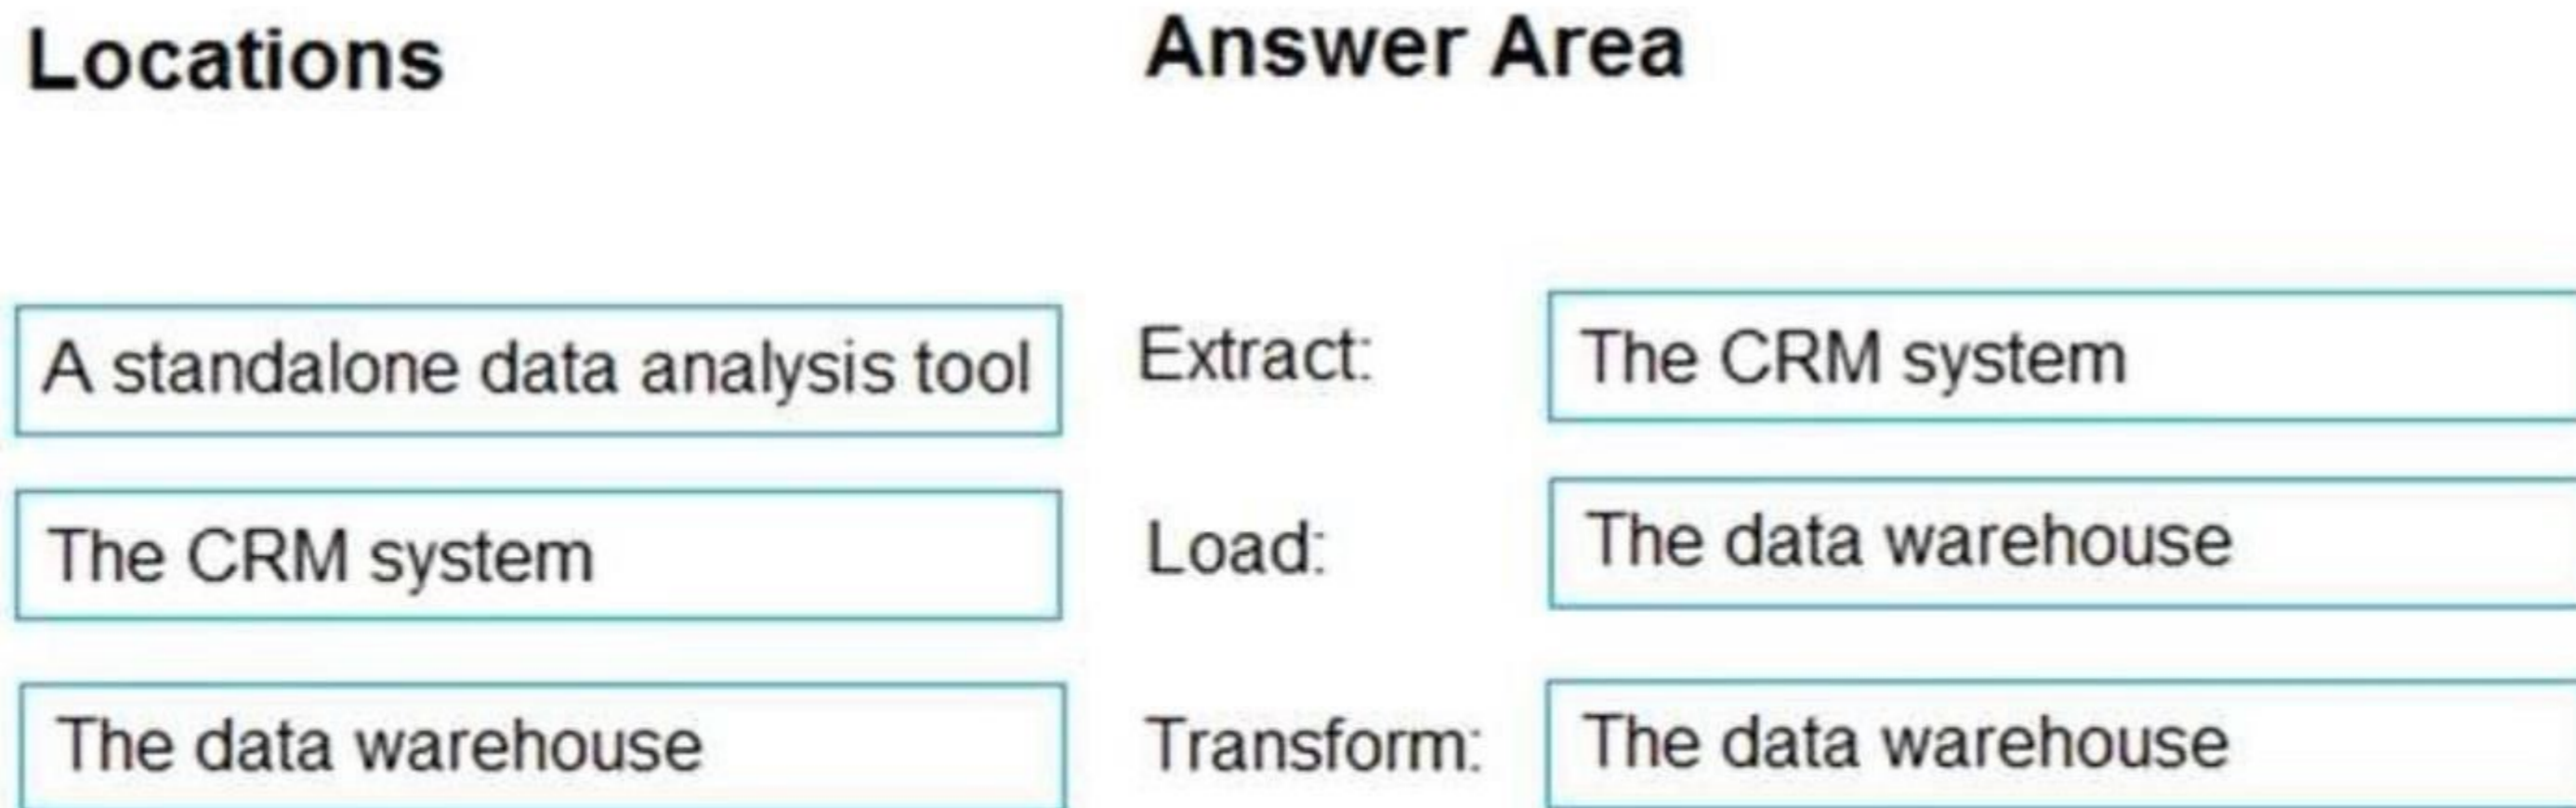
NEW QUESTION: 83

- Which of the following is a Microsoft SQL Server data type?
- A. BIT
 - B. VARCHAR
 - C. NVARCHAR
 - D. CHAR

□□□□: □□ □□□ 1□□□□.



Answer:



Explanation:
Extract, load, and transform (ELT) differs from ETL solely in where the transformation takes place. In the ELT pipeline, the transformation occurs in the target data store. Instead of using a separate transformation engine, the processing capabilities of the target data store are used to transform data (...) In practice, the TARGET data store is a DATA WAREHOUSE using either a Hadoop cluster (using Hive or Spark) or a SQL dedicated pools on Azure Synapse Analytics.

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl> = DW for the Transform.

□□□ □□
□□ □ □□□ □□ □□□ □□□□□ □□ □□□□□. □□□ □□□ □□□□ □□□□□.
□□□□: □□ □□□ 1□□□□.

Answer Area



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

Answer:

Answer Area

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

Explanation:
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: Yes

Batch processing can have output to SQL Database, Hive, HBase, Spark SQL.

Box 3: Yes

Reference:

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

NEW QUESTION: 87

Which of the following is a declarative big data processing language used in many Hadoop distributions, including HDInsight. It is particularly useful for processing data that is unstructured or semi-structured.

- A. Pig
- B. Hive
- C. Tez
- D. MapReduce

Answer: C (LEAVE A REPLY)

Pig is a declarative big data processing language used in many Hadoop distributions, including HDInsight. It is particularly useful for processing data that is unstructured or semi-structured.

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

NEW QUESTION: 88

Core(SQL) API is available for Azure Cosmos DB. Which of the following is not a supported API?

Which of the following is not a supported API? (200 characters) A: REST API B: .NET SDK C: Java SDK D: API

- A. REST API
- B. .NET SDK
- C. Java SDK
- D. API

Answer: A,C (LEAVE A REPLY)

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

NEW QUESTION: 89

The pipeline engine used in Azure Data Factory is also available in which of the following Azure services?

Answer Area

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

	▼
<input type="checkbox"/>	Azure Databricks.
<input type="checkbox"/>	Azure HDInsight.
<input type="checkbox"/>	Azure SQL Database.
<input type="checkbox"/>	Azure Synapse Analytics.

Answer:

Answer Area



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

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

NEW QUESTION: 90

□□□ □□
□□ □ □□□ □□ □□□□ □□ □□□□□. □□□ □□□ □□□□ □□□□□.
□□□□: □□ □□□ 1□□□□.

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"/>

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"/>

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

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

NEW QUESTION: 91

□□ □□□□ □□□ □□ □□ □□□□□□□ □□□□.

ProductName	Quantity
Product1	100
Product2	129
Product3	176

Product1 □ □□ □□□ 270 □□ □□□□□ SQL □□□□ □□ □□ □□□□ □□□?

- A. □□
- B. □□
- C. □□□□
- D. □□

Answer: C (LEAVE A REPLY)

Update statement is used to change the value of a column.

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

DP-900-KR □□ □□□ □□□□□ □□ DumpTop □□ □□□□ □□□ DP-900-KR □□! DumpTop □ □□ **DP-900-KR** □□ □□□ □□□□□□, DumpTop DP-900-KR □□ □□□ □□□□□□□□ □□□ □□□□ □□□. □□□□ □□□ □□□□ □□ DumpTop DP-900-KR □□□ □□□□□. <https://www.dumptop.com/Microsoft/DP-900-KR-dump.html> (316 Q&As Dumps, **30%OFF Special Discount: KrDump**)

NEW QUESTION: 92

□□□ □□
□□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

Answer Area

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

natively support the analysis of relationships between entities.



Answer:

Answer Area

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

natively support the analysis of relationships between entities.

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

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

NEW QUESTION: 93

□□□ □□
□□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

Answer Area

A key/value data store is optimized for

▼
enforcing constraints
simple lookups
table joins
transactions



Answer:

Answer Area

A key/value data store is optimized for

	▼
enforcing constraints	
simple lookups	
table joins	
transactions	

Explanation:

Box 1: simple lookups

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

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

Reference:

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

NEW QUESTION: 94

Which of the following is a feature of Azure Data Factory?

- A. Data integration
- B. Data orchestration
- C. Data migration
- D. Data replication

Answer: B (LEAVE A REPLY)

Trigger.

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

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

NEW QUESTION: 95

Which of the following is a feature of Azure Database for MariaDB?

- A. High availability
- B. Automatic backup and restore
- C. Point-in-time recovery
- D. Read replicas

Answer: B (LEAVE A REPLY)

Azure Database for MariaDB is a relational database service with the following features:

- * Build-in high availability without any extra cost
- * Autoscaling whenever required, within seconds
- * Enterprise-level security and compliance
- * Predictable performance with the pay-as-you-go model
- * Auto backups and point-in-time-restore for up to 35 days
- * Protection and security of confidential data in motion as well as at rest Azure Database for MariaDB has a limitation that the mysql system database in Azure Database for MariaDB is a read-only as it supports a number of PaaS functionality. You can't make any changes to the mysql system database.

Option A is incorrect. Azure Database for MariaDB delivers build-in high availability without any additional cost and predictable performance with a pay-as-you-go pricing model.

Option B is correct. The mysql system database in Azure Database for MariaDB is a read-only as it supports a number of PaaS functionality. You can't make any changes to the mysql system database.

Option C is incorrect. Azure Database for MariaDB delivers autoscaling in seconds when required, auto backups, and power-in-time-restore.

Option D is incorrect. Azure Database for MariaDB delivers secured protection of sensitive data in motion as well as at rest, and also offer enterprise-level compliance and security.

References: To know more about the Azure Database for MariaDB, refer to the links below:

<https://docs.microsoft.com/en-us/azure/mariadb/overview>

<https://docs.microsoft.com/en-us/azure/mariadb/concepts-limits>

NEW QUESTION: 96

□□□ □□
□□□ □□□□ □□□□ □□ □□□□□.

Answer Area

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

ALTER
JOIN
SET
WHERE

Answer:

Answer Area

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

ALTER
JOIN
SET
WHERE

NEW QUESTION: 97

□□□ □□
□□□ □□□□ □□□□ □□ □□□□□.

Answer Area

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

▼ data.

relational
semi-structured
structured
unstructured

Answer:

Answer Area

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

▼ data.

relational
semi-structured
structured
unstructured




NEW QUESTION: 98

□□□ □□
□□□ □□□□ □□□□ □□ □□□□□.

Answer Area

▼ 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:

Answer Area

▼

- A data analyst
- A data engineer
- A data scientist

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



Explanation:

A data analyst is responsible for identifying which business rules must be applied to the data of a company. Their role involves understanding business needs, analyzing data, and ensuring that appropriate business rules are applied to support decision-making and reporting.

NEW QUESTION: 99

QUESTION 99: A data analyst is responsible for identifying which business rules must be applied to the data of a company. Their role involves understanding business needs, analyzing data, and ensuring that appropriate business rules are applied to support decision-making and reporting.

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:

Analytics Types

Answer Area

Microsoft

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

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 1 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/>

NEW QUESTION: 100

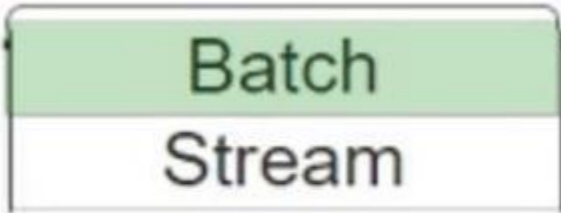
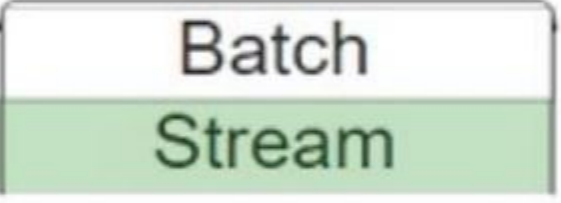
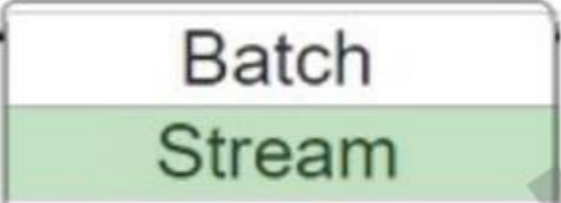
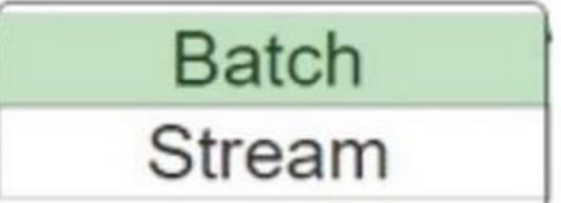
□□□ □□

□ □□□□□□ □□□ □□ □□ □□□ □□□ □□ □ □□ □□ □□□□ □□□□□.

□□□□□ □□□□ □□□□ □ □□□□□□□ □□ □□ □□ □□ □□□□ □□□□□.

Processing type	Application
Batch Stream	Employee payroll processing and generating payroll checks
Batch Stream	Reporting the number of users and bandwidth user for an online game
Batch Stream	Identifying defected manufacturing errors to automatically reject failing parts
Batch Stream	Setting inventory stocking levels based on seasonal sales volume

Answer:

Processing type	Application
	Employee payroll processing and generating payroll checks
	Reporting the number of users and bandwidth user for an online game
	Identifying defected manufacturing errors to automatically reject failing parts
	Setting inventory stocking levels based on seasonal sales volume

Explanation:

The following two applications are best suited to batch processing:

- Employee payroll processing and generating payroll checks
- Setting inventory stocking levels based on seasonal sales volume

Data for batch processing is collected over time, often from different data sources, and it is processed as a dataset that includes a range of rows or all rows in the dataset. Batch processing is designed to handle processing of large datasets. There is typically a long latency between data collection and data processing.

The following two applications are best suited to stream processing:

Option B is incorrect. Follow setup for SDK logging is a checkpoint in the checklist for troubleshooting issues when you are working with Azure Cosmos DB .NET SDK.

Option C is incorrect. Log the SQL Query Metrics from all the query responses is a checkpoint in the checklist for troubleshooting issues when you are working with Azure Cosmos DB .NET SDK.

Option D is correct. Disable the SDK logging is not a checkpoint in the checklist for troubleshooting issues when you are working with Azure Cosmos DB .NET SDK. The right one is - Enable the SDK logging.


Reference: To know more about the troubleshooting in Azure Cosmos DB .NET SDK, refer to the link below.

<https://docs.microsoft.com/en-us/azure/cosmos-db/troubleshoot-dot-net-sdk>

NEW QUESTION: 105

□□□ □□

□□□ □□□□ □□□□ □□ □□□□□.



Answer:



Explanation:

Azure Synapse Analytics is an platform as a service (PAAS) that combines data integration, warehousing, and analytics into one solution. Azure Synapse Analytics offers cloud-based, relational data warehousing services, massively parallel processing (MPP) scale-out technology, and enough computational power to efficiently manage petabytes and petabytes of data.

Incorrect:

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

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

Reference:

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

<https://www.integrate.io/blog/what-is-azure-synapse-analytics>

NEW QUESTION: 106

Answer Area

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

Explanation:

Box 1: Yes

Throughput is the number of transactions per sec. Normalization decrease the amount of data that needs to be written, hence increasing the amount or transactions that can be written per sec.

Box 2: No

Analytics systems are denormalized to increase performance.

Transactional database systems are normalized to increase data consistency.

Box 3: Yes

Transactional database systems are more normalized and requires more joins.

Reference:

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

NEW QUESTION: 108

□□□ □□

□□□ □□□□ □□□□ □□ □□□□□.

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.

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

Answer:

Answer Area

Microsoft

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

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

NEW QUESTION: 109

□□□ □□
 □□□ □□□□□ □□□□□ □□□ □□□ □□□□□.

Answer Area

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

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

Answer:

Answer Area Microsoft

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:

Distributions map to Compute nodes for processing. As you pay for more compute resources, distributions are remapped to available Compute nodes.
<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/massively-parallel-processing-mpp-architecture> [see the architecture diagram]

NEW QUESTION: 110


□□□ □□
 □□ □ □□□ □□ □□□ □□□□□ □□ □□□□□. □□□ □□□ □□□□ □□□□□.
 □□□□: □□ □□□ 1□□□□.

Answer Area

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

Answer:

Answer Area

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

Explanation:

Box 1: No

The API determines the type of account to create. Azure Cosmos DB provides five APIs: Core (SQL) and MongoDB for document data, Gremlin for graph data, Azure Table, and Cassandra. Currently, you must create a separate account for each API.

Box 2: Yes

Azure Cosmos DB uses partitioning to scale individual containers in a database to meet the performance needs of your application. In partitioning, the items in a container are divided into distinct subsets called logical partitions. Logical partitions are formed based on the value of a partition key that is associated with each item in a container.

Box 3: No

Logical partitions are formed based on the value of a partition key that is associated with each item in a container.

Reference:

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

NEW QUESTION: 111

Which cloud service model is used to host a virtual machine in Azure? The client manages the operating system, applications, and data, while the cloud provider manages the hardware and network infrastructure.

- A. Infrastructure as a Service (IaaS)
- B. Software as a Service (SaaS)
- C. Platform as a Service (PaaS)
- D. Serverless Compute Service

Answer: (SHOW ANSWER)

SaaS (Software as a Service): This is a fully managed cloud service where the client uses the software without worrying about infrastructure, platform, or application management. The cloud provider handles everything, including updates, security, and maintenance.

NEW QUESTION: 112

Which tool is used to create and manage Azure SQL Server Notebooks? The tool allows you to create and manage notebooks in the Azure portal, and it provides a rich text editor and a query editor.

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

Answer: B (LEAVE A REPLY)

Azure Data Studio. It allows you to save your query and even create charts or simple reports that you can re-visit. <https://www.mssqltips.com/sqlservertip/5997/create-sql-server-notebooks-in-azure-data-studio/>

NEW QUESTION: 113

Which Azure service is used to store and manage metadata for data lakes and data warehouses? The service provides a scalable and secure metadata store that can be accessed from any location.

Once you select a tile, you will move to the report page where that visualization was created.

Option A is incorrect. Not visualizations but the Power BI dashboard is known as Canvas.

Option B is incorrect. There is nothing like Elements on the Power BI dashboard.

Option C is correct. The visualizations that are present on the Power BI Dashboard are known as Tiles.

Option D is incorrect. The visualizations come from the reports, these are not the reports.

Reference: To know more about the Visualization in Microsoft Power BI, refer to the documentation below from Azure:

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

NEW QUESTION: 116

□□□ □ □□ □□

□□ □□□ □□□ □□□ □□ □□□□□□.

□□□□□ □□ □□□ □□ □□□ □□□ □□□□ □□□ □□□□□. □ □□□ □ □, □ □ □□ □□ □□ □□□□ □□ □ □□□□.

□□□□: □□□ □□ □□□ 1□□□□.

Job role

Data analyst

Data engineer

Database administrator

Answer Area

Assign permissions to users.

Derive meaning from data by creating visualizations and charts.

Implement which pipelines to transfer and transform data between systems.

Answer:

Job role	Answer Area
Database administrator	Assign permissions to users.
Data analyst	Derive meaning from data by creating visualizations and charts.
Data engineer	Implement which pipelines to transfer and transform data between systems.

NEW QUESTION: 117

□□□ □ □□ □□

Azure □□□ □□□ □□ □□□ □□□ □□□ □□□□□□.

□□□□□ □□ □□□ □□ □□ □□□ □□□ □□□□□□□□. □ □□ □□□ □ □, □ □ □□ □□ □□ □□□□ □□ □ □□□□.

□□□□: □□□ □□ □□□ 1□□□□.

Components

- Dataset
- Linked service
- Mapping data flow
- Pipeline

Answer Area

- Component
- Component
- Component

A representation of data structures within data stores

The information used to connect to external resources

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



Answer:

The screenshot shows the 'Answer Area' with three boxes selected: 'Dataset', 'Linked service', and 'Pipeline'. Each box is accompanied by its definition: 'Dataset' is 'A representation of data structures within data stores', 'Linked service' is 'The information used to connect to external resources', and 'Pipeline' is 'A logical grouping of activities that performs a unit of work and can be scheduled'. The 'Components' list on the left remains the same as in the top image.

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

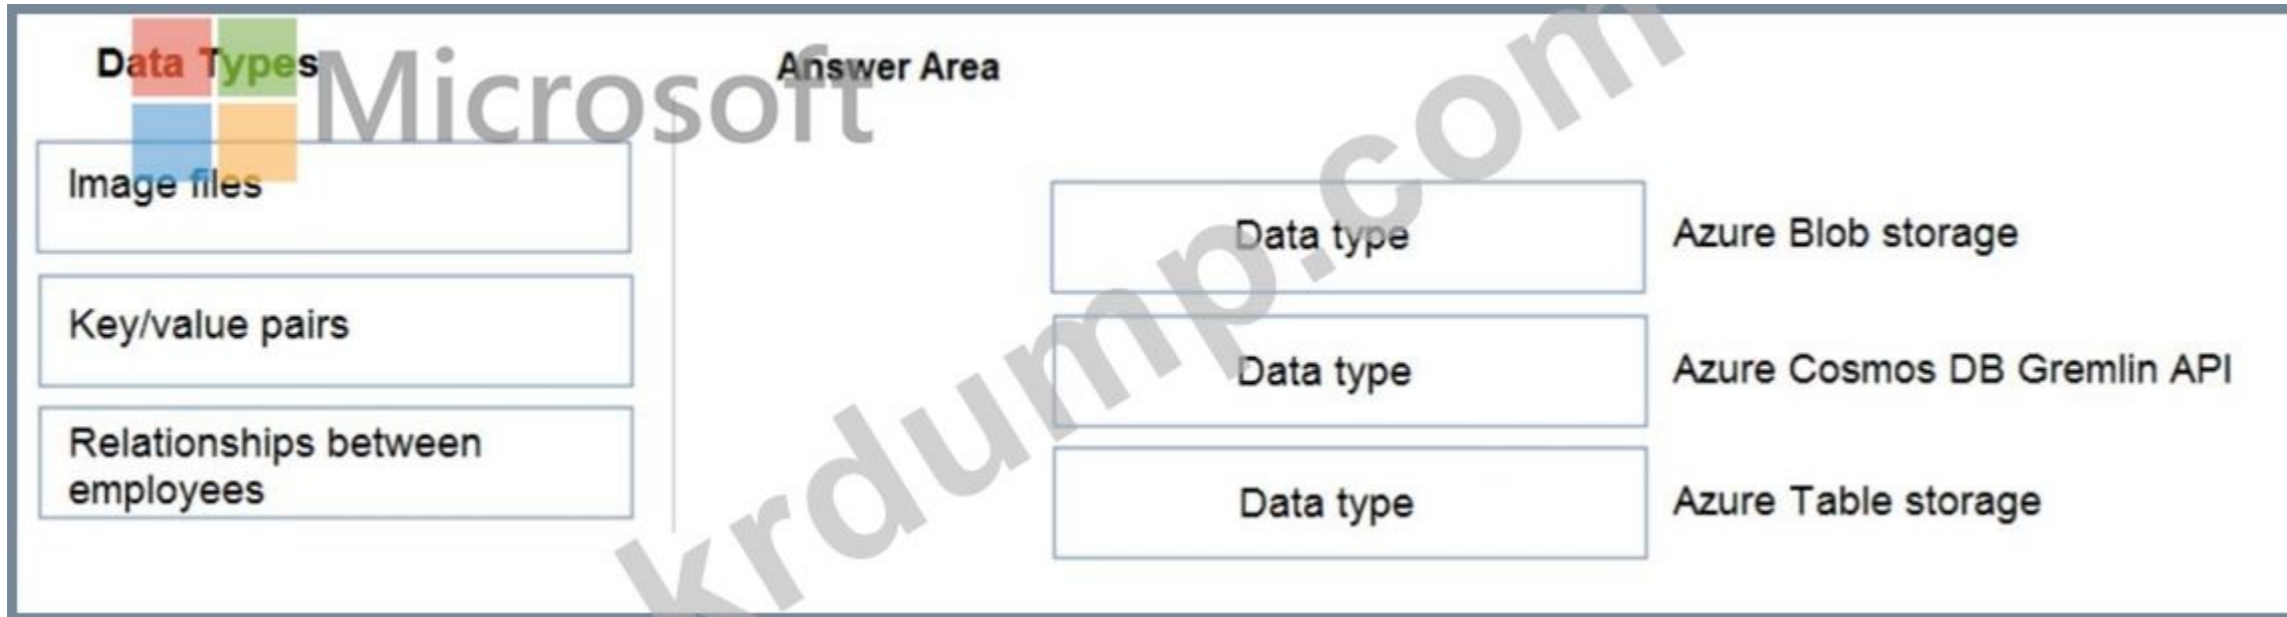
A pipeline is a logical grouping of activities that together perform a task.

Reference:

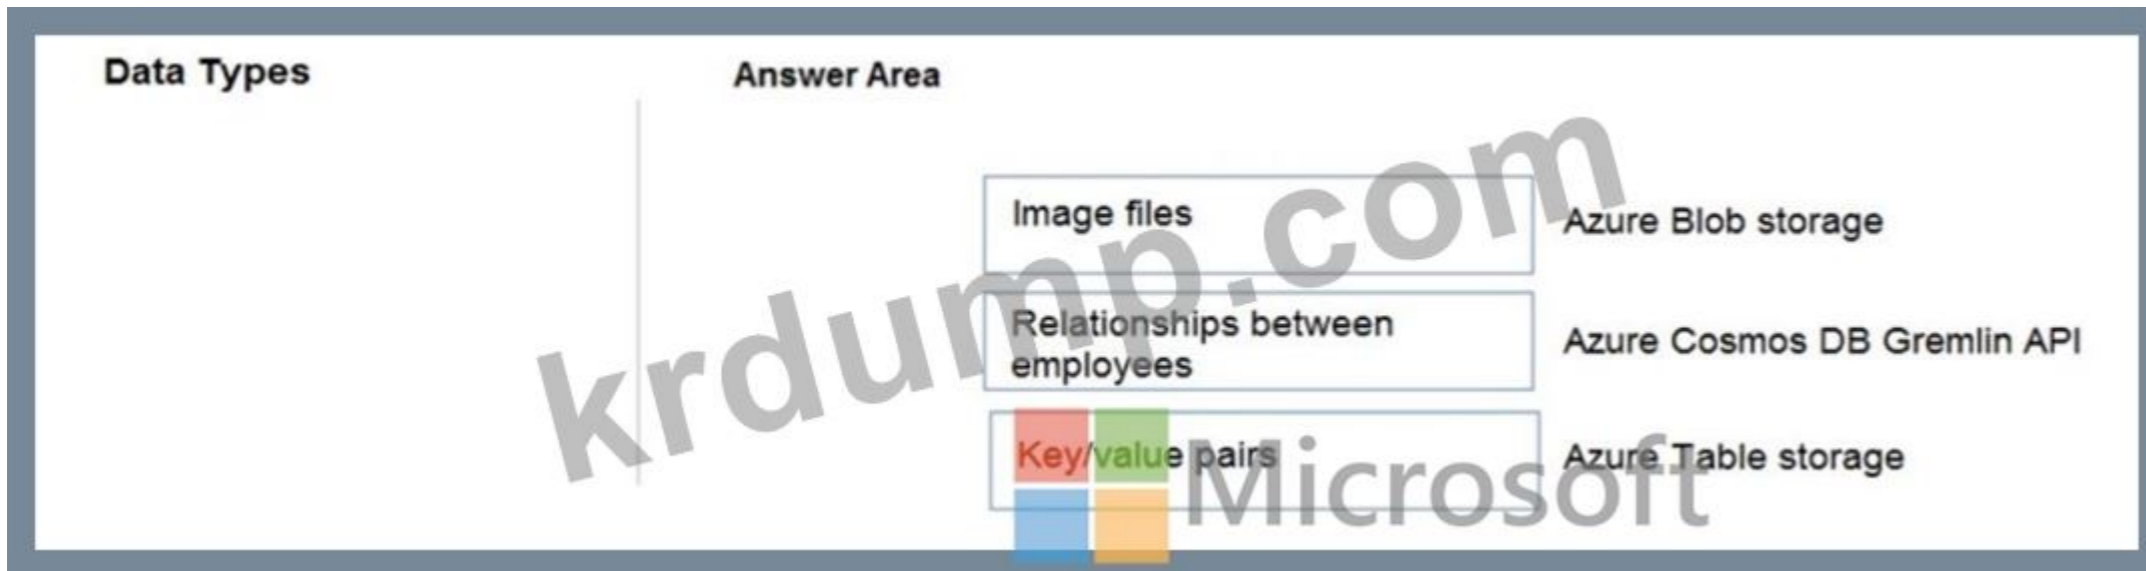
<https://k21academy.com/microsoft-azure/dp-100/datastores-and-datasets-in-azure/>

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

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



Answer:



Explanation:

Box 1: Image files

Image files are what you upload on a blob, as they are unstructured data.

Box 2: Relationship between employees

Relationships between employees are context based, meaning the entities will exist as nodes in a graph and you will use the Cosmos DB Gremlin API.

Box 3: Key/value pairs

Key/value data can be stored in a Cosmos DB, but it can also be stored in an Azure Table.

Reference:

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