



Microsoft

Exam Questions DP-900

Microsoft Azure Data Fundamentals

About ExamBible

[Your Partner of IT Exam](#)

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

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

- A. Azure SQL Database Hyperscale
- B. Azure SQL managed instance
- C. Azure SQL Database serveless
- D. Azure SQL Database elastic pod

Answer: B

NEW QUESTION 2

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

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

Answer: A

Explanation:

Power BI Report Builder is the standalone tool for authoring paginated reports for the Power BI service. <https://docs.microsoft.com/en-us/power-bi/paginated-reports/paginated-reports-quickstart-aw>

NEW QUESTION 3

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

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

Answer: B

NEW QUESTION 4

Select the answer that correctly completes the sentence.

Answer Area

Backing up and restoring data are the responsibilities of a

- data engineer.
- data analyst.
- data engineer.**
- data scientist.
- database administrator.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Backing up and restoring data are the responsibilities of a

- data engineer.
- data analyst.
- data engineer.**
- data scientist.
- database administrator.

NEW QUESTION 5

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

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

NOTE: Each correct selection is worth one point.

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

Answer: AC

NEW QUESTION 6

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

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

Which type of data store should you use?

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

Answer: D

NEW QUESTION 7

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

Answer Area

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

cognitive

descriptive

predictive

prescriptive

- A. Mastered
- B. Not Mastered

Answer: A

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?

NEW QUESTION 8

Which three actions are used when performing diagnostic analytics?

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

Answer: A

NEW QUESTION 9

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

Relational data is stored in

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

NEW QUESTION 10

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

Transparent Data Encryption (TDE) encrypts

a column to protect data at rest and in transit.
queries and their results in order to protect data in transit.
the database to protect data at rest.
the server to protect data at rest.

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Transparent Data Encryption (TDE) encrypts

a column to protect data at rest and in transit.
queries and their results in order to protect data in transit.
the database to protect data at rest.
the server to protect data at rest.

NEW QUESTION 10

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

In a data warehousing workload, data

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

- A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 15

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

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

Answer: C

NEW QUESTION 18

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 22

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

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

Box 2: Yes

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



Box 3: No

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

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

NEW QUESTION 24

What is a characteristic of a non-relational database?

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

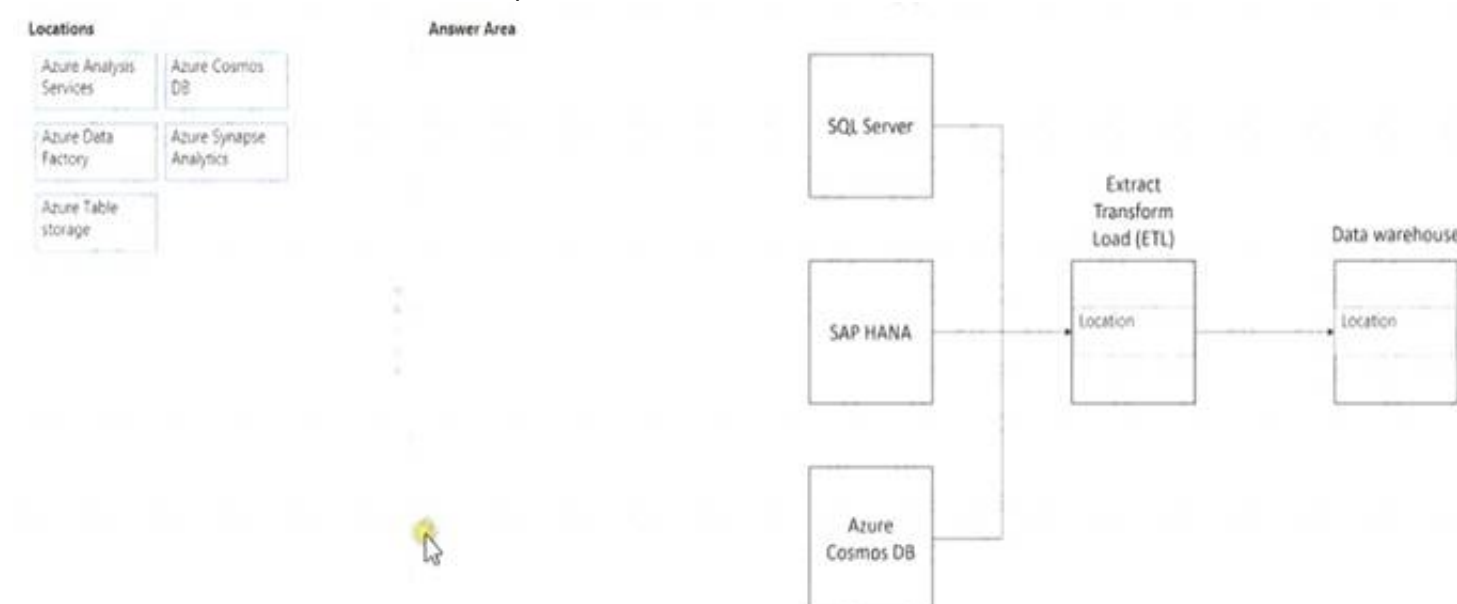
Answer: C

NEW QUESTION 27

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.

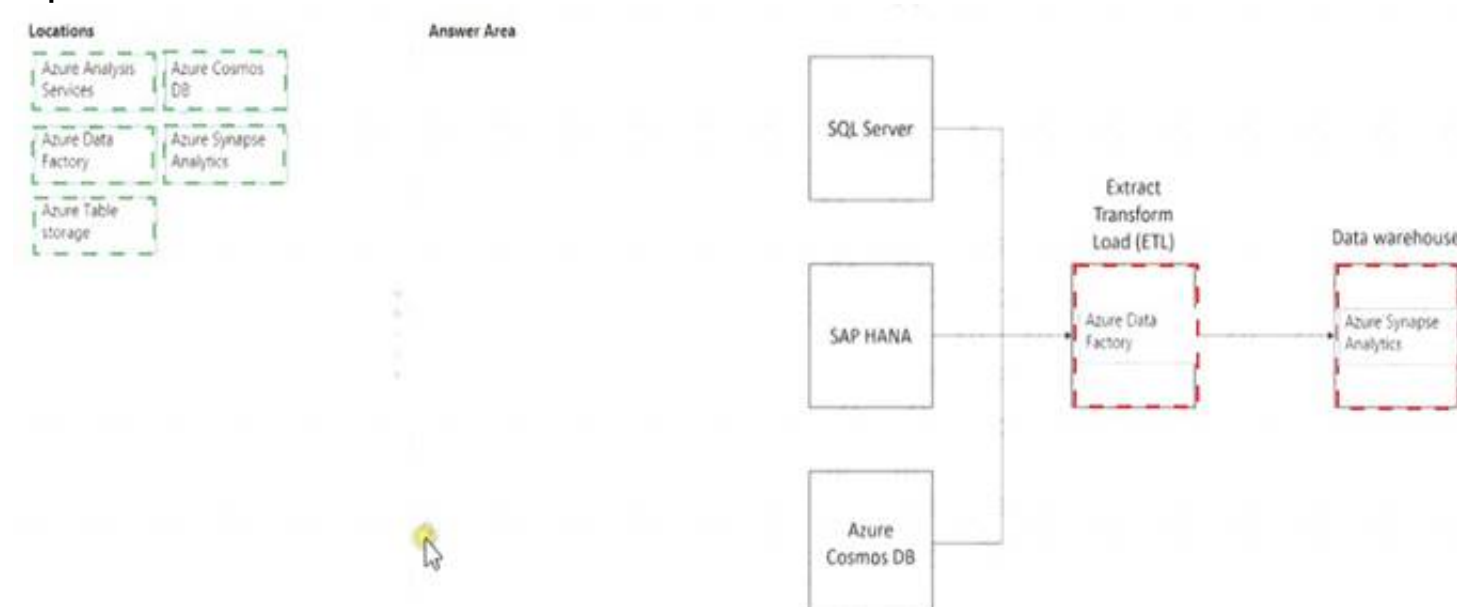
NOTE: Each correct match is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

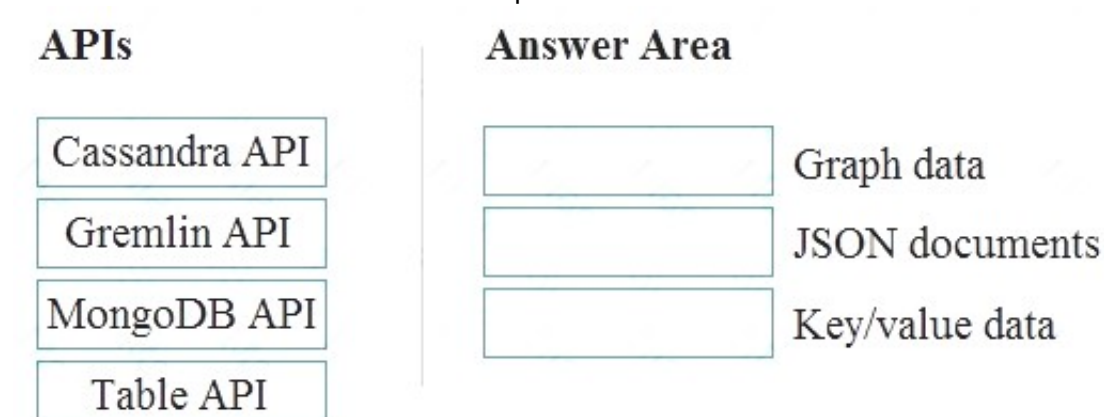


NEW QUESTION 32

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

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

NOTE: Each correct match is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

APIs

Cassandra API

Gremlin API

MongoDB API

Table API

Answer Area

Gremlin API

Graph data

MongoDB API

JSON documents

Table API

Key/value data

NEW QUESTION 34

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

▼

A heap

A stored procedure

A view

An index

presents content defined by a query.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

A View

NEW QUESTION 38

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

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

Answer: B

NEW QUESTION 39

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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 checked="" 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"/>

NEW QUESTION 43

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

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

Answer: D

NEW QUESTION 44

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 45

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

<div><div></div><div>A clustered index</div><div>A FileTable</div><div>A foreign key</div><div>A stored procedure</div></div>	is an object associated with a table that sorts and stores the data rows in the table based on their key values.
-------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

<div><div></div><div>A clustered index</div><div>A FileTable</div><div>A foreign key</div><div>A stored procedure</div></div>	is an object associated with a table that sorts and stores the data rows in the table based on their key values.
-------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------

NEW QUESTION 50

Which two Azure services can be used to provision Apache Spark clusters? Each correct answer presents a complete solution. (Choose two.)
NOTE: Each correct selection is worth one point.

- A. Azure Time Series Insights
- B. Azure HDInsight

- C. Azure Databricks
- D. Azure Log Analytics

Answer: BC

NEW QUESTION 53

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

To configure an Azure Storage account to support access control lists

that have object-level permissions,

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Table Description automatically generated

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.

NEW QUESTION 55

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

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

Answer: B

NEW QUESTION 59

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.

Components

Dataset

Linked service

Mapping data flow

Pipeline

Answer Area

Component

A representation of data structures within data stores

Component

The information used to connect to external resources

Component

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

- A. Mastered
- B. Not Mastered

Answer: A

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.

NEW QUESTION 62

Select the answer that correctly completes the sentence.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 63

Which setting can only be configured during the creation of an Azure Cosmos DB account?

- A. geo-redundancy
- B. multi-region writes
- C. production or non-reduction account type
- D. API

Answer: D

Explanation:

You must create a separate account for each API.

NEW QUESTION 64

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

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

<input type="text"/>	▼
distributes processing across compute nodes.	
distributes processing across control nodes.	
redirects client connections across compute nodes.	
redirects client connections across control nodes.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

<input type="text"/>	▼
distributes processing across compute nodes. !	
distributes processing across control nodes.	
redirects client connections across compute nodes.	
redirects client connections across control nodes.	

NEW QUESTION 67

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
A pipeline is a representation of a data structure within Azure Data Factory	<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"/>

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Box 1: No
A pipeline is a logical grouping of activities that together perform a task.
Box 2: Yes
You can construct pipeline hierarchies with data factory.
Box 3: Yes
A pipeline is a logical grouping of activities that together perform a task.

NEW QUESTION 71

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

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

Answer: D

NEW QUESTION 74

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

A block of code that runs in a database is called

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated with medium confidence

NEW QUESTION 75

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

Batch workloads

process data in memory, row-by-row.
collect and process data at most once a day.
process data as new data is received in near real-time.
collect data and then process the data when a condition is met.

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Batch workloads

process data in memory, row-by-row.	
collect and process data at most once a day.	
process data as new data is received in near real-time.	
collect data and then process the data when a condition is met.	

NEW QUESTION 79

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

12345

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"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

12345

Answer Area

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

NEW QUESTION 82

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

Statements	Yes	No
You must apply patches to Azure SQL databases regularly.	<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"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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"/>

NEW QUESTION 84

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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"/>

NEW QUESTION 89

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

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

Answer: B

NEW QUESTION 90

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.

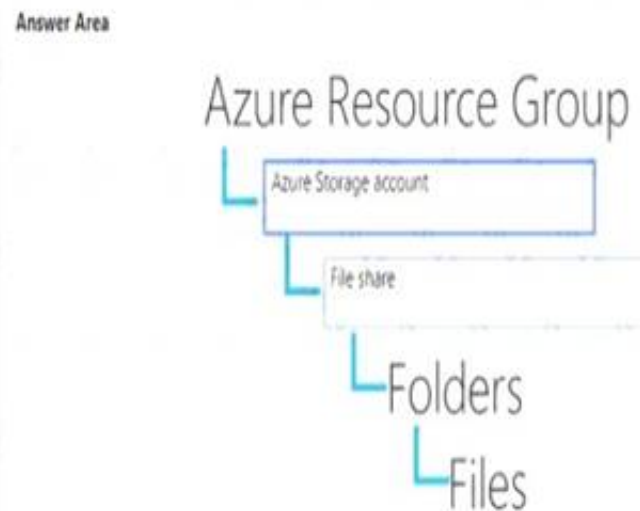
NOTE: Each correct match is worth one point.



- A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 95

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

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

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated with medium confidence

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

NEW QUESTION 98

You need to create an Azure Storage account.

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

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

NOTE: Each correct selection is worth one pint.

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

Answer: AC

Explanation:

_ Azure Storage offers two options for copying your data to a secondary region: Geo-redundant storage (GRS)

_ Geo-zone-redundant storage (GZRS)

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

NEW QUESTION 99

Select the answer that correctly completes the sentence. See the answer below:

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

An Azure Storage account must have

Hierarchical namespace

enabled to support Azure Data Lake Storage.

NEW QUESTION 102

Match the security components to the appropriate scenarios.

To answer, drag the appropriate component from the column on the left to its scenario on the right. Each component may be used once, more than once, or not at all. NOTE: Each correct match is worth one point.

Components

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.

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

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.

NEW QUESTION 104

You need to recommend a non-relational data store that is optimized for storing and retrieving files, videos, audio stream, and virtual disk images. The data store must store data, some metadata, and a unique ID for each file. Which type of data store should you recommend?

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

Answer: C

Explanation:

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

NEW QUESTION 107

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

Answer Area

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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).
```

NEW QUESTION 112

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

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

Answer: D

NEW QUESTION 113

You have an Azure SQL database that you access directly from the internet. You recently changed your external IP address. After changing the IP address, you can no longer access the database. You can connect to other resources in Azure. What is a possible cause of the issue?

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

Answer: A

Explanation:

The Azure SQL Database firewall lets you decide which IP addresses may or may not have access to either your Azure SQL Server or your Azure SQL database. When creating an Azure SQL Database, the firewall needs to be configured before anyone will be able to access the database. By default, no external access to your SQL Database will be allowed until you explicitly assign permission by creating a firewall rule.

NEW QUESTION 118

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

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

Answer: B

NEW QUESTION 121

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

Transcribing audio files is an example of

cognitive

descriptive

predictive

prescriptive

 analytics.

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Transcribing audio files is an example of

cognitive

descriptive

predictive

prescriptive

 analytics.

NEW QUESTION 125

Your company needs to design a database that illustrates the relationships between utilization levels of individual network devices across a local area network. Which type of data store should you use?

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

Answer: A

Explanation:

Data as it appears in the real world is naturally connected. Traditional data modeling focuses on defining entities separately and computing their relationships at runtime. While this model has its advantages, highly connected data can be challenging to manage under its constraints. A graph database approach relies on persisting relationships in the storage layer instead, which leads to highly efficient graph retrieval operations. Azure Cosmos DB's Gremlin API supports the property graph model.

NEW QUESTION 130

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

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

Answer: A

NEW QUESTION 132

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

Answer Area		
Statements	Yes	No
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"/>

- A. Mastered
B. Not Mastered

Answer: A

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.

NEW QUESTION 137

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

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

Answer: C

NEW QUESTION 138

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

natively support the analysis of relationships between entities.

Column family databases

Document databases

Graph databases

Key-value stores

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

NEW QUESTION 140

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

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

account

container

database

item

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Table Description automatically generated

NEW QUESTION 144

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

Answer Area			
		Statements	
			YesNo
		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"/>

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

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"/>

NEW QUESTION 149

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

Statements	Yes	No
Normalization involves eliminating relationships between database tables.	<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"/>

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

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"/>

NEW QUESTION 152

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.

Language	Answer Area
JavaScript	<input type="text"/> Azure Data Explorer
KQL	<input type="text"/> Azure Stream Analytics
Python	<input type="text"/> Azure Synapse Analytics
SQL	<input type="text"/>

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Language	Answer Area
JavaScript	<input type="text" value="KQL"/> Azure Data Explorer
KQL	<input type="text" value="SQL"/> Azure Stream Analytics
Python	<input type="text" value="Python"/> Azure Synapse Analytics
SQL	

NEW QUESTION 157

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

Statements	Yes	No
Azure 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"/>

- A. Mastered
- B. Not Mastered

Answer: A

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.

NEW QUESTION 159

What are three characteristics of an Online Transaction Processing (OLTP) workload? Each correct answer presents a complete solution. (Choose three.)
NOTE: Each correct selection is worth one point.

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

Answer: BDF

NEW QUESTION 162

A team of developers has computers that run Windows 10 and Ubuntu Desktop.

The developers need to connect to and query an Azure SQL database from the computers. The developers require code assistance features such as IntelliSense. What should the developers use ?

- A. Azure Data Studio
- B. sq1cmd
- C. Microsoft SQL Server Management Studio (SSMS)
- D. Azure Data Explorer

Answer: A

Explanation:

Azure Data Studio is a cross-platform database tool for data professionals who use on-premises and cloud data platforms on Windows, macOS, and Linux. Azure Data Studio offers a modern editor experience with IntelliSense, code snippets, source control integration, and an integrated terminal.

NEW QUESTION 165

Which Azure SQL offering provides near-100 percent compatibility with on-premises Microsoft SQL Server instances, while providing automated updates,

backups, and maintenance tasks?

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

Answer: A

NEW QUESTION 167

Select the answer that correctly completes the sentence.

In a SQL statement,

FROM and WHERE

GROUP BY and ORDER BY

JOIN and MERGE

MAX and MIN

 are examples of aggregate functions.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Table Description automatically generated

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

NEW QUESTION 168

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

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

Answer: D

NEW QUESTION 171

Which type of database is Azure Database for PostgreSQL?

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

Answer: A

NEW QUESTION 173

Select the answer that correctly completes the sentence.

Answer Area

A clustered index

A view

A clustered index

A stored procedure

A nonclustered index

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A clustered index

A view

A clustered index

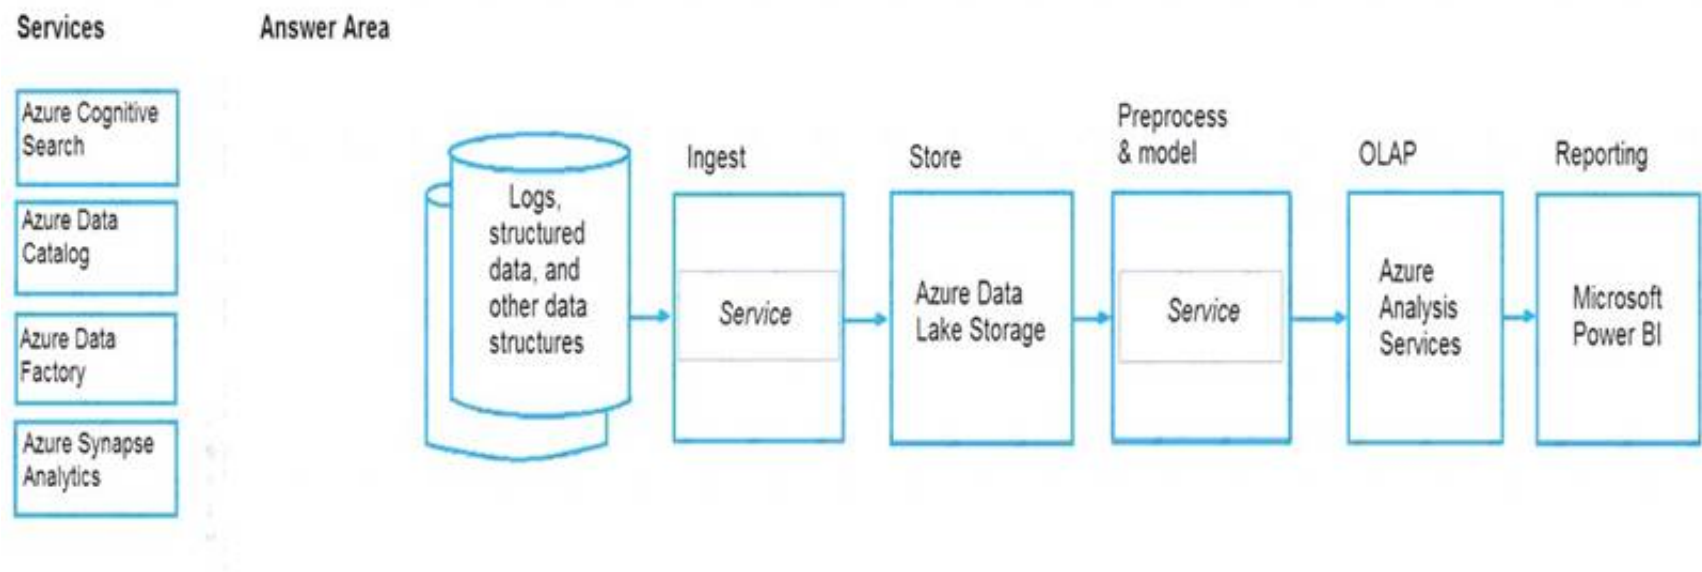
A stored procedure

A nonclustered index

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

NEW QUESTION 175

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.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box Ingest: Azure Data Factory
You can build a data ingestion pipeline with Azure Data Factory (ADF). Box Preprocess & model: Azure Synapse Analytics
Use Azure Synapse Analytics to preprocess data and deploy machine learning models.

NEW QUESTION 178

Your company needs to design a database that shows how changes traffic in one area of a network affect other components on the network.
Which type of data store should you use?

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

Answer: B

Explanation:

Data as it appears in the real world is naturally connected. Traditional data modeling focuses on defining entities separately and computing their relationships at runtime. While this model has its advantages, highly connected data can be challenging to manage under its constraints. A graph database approach relies on persisting relationships in the storage layer instead, which leads to highly efficient graph retrieval operations. Azure Cosmos DB's Gremlin API supports the property graph model.

NEW QUESTION 183

Select the answer that correctly completes the sentence.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 184

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

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

Answer: A

NEW QUESTION 188

Select the answer that correctly completes the sentence.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 193

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

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

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"/>

NEW QUESTION 197

When can you use an Azure Resource Manager template?

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

Answer: A

Explanation:

You can automate deployments and use the practice of infrastructure as code. In code, you define the infrastructure that needs to be deployed. To implement infrastructure as code for your Azure solutions, use Azure Resource Manager templates (ARM templates). The template is a JavaScript Object Notation (JSON) file that defines the infrastructure and configuration for your project. The template uses declarative syntax, which lets you state what you intend to deploy without having to write the sequence of programming commands to create it. In the template, you specify the resources to deploy and the properties for those resources.

NEW QUESTION 200

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

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.

NEW QUESTION 204

You need to perform hybrid transactional and analytical processing (HTAP) queries against Azure Cosmos DB data sources by using Azure Synapse Analytics. What should you use?

- A. Synapse pipelines
- B. a Synapse SQL pool
- C. Synapse Link
- D. Synapse Studio

Answer: C

NEW QUESTION 205

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 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.

VALUES

CREATE

GROUP BY

ORDER BY

SELECT

ANSWER AREA

Value

SalesDate, AVG(SalesAmount)

FROM Sales

Value

SalesDate

- A. Mastered
- B. Not Mastered

Answer: A

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;
```

NEW QUESTION 209

Select the answer that correctly completes the sentence.

Answer Area

An Azure Storage account must have

Blob soft delete
Hierarchical namespace
Large file shares
Versioning

enabled to support Azure Data Lake Storage.

- A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 210

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

Answer Area

Statements	Yes	No
Azure SQL managed instance supports cross-database queries.	<input type="radio"/>	<input type="radio"/>
Azure SQL managed instance supports user-created backups.	<input type="radio"/>	<input type="radio"/>
Azure SQL managed instance databases can be restored to SQL Server on an Azure virtual machine.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Azure SQL managed instance supports cross-database queries.	<input checked="" type="radio"/>	<input type="radio"/>
Azure SQL managed instance supports user-created backups.	<input type="radio"/>	<input checked="" type="radio"/>
Azure SQL managed instance databases can be restored to SQL Server on an Azure virtual machine.	<input checked="" type="radio"/>	<input type="radio"/>

NEW QUESTION 214

What is the primary purpose of a data warehouse?

- A. to provide answers to complex queries that rely on data from multiple sources
B. to provide transformation services between source and target data stores
C. to provide read-only storage of relational and non-relational historical data
D. to provide storage for transactional line-of-business (LOB) applications

Answer: C

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

NEW QUESTION 216

Select the answer that correctly completes the sentence.

Answer Area

A view

A heap

A stored procedure

A view

An index

presents content defined by a query.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A view

A heap

A stored procedure

A view

An index

presents content defined by a query.

NEW QUESTION 220

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Box 1: Yes

For read access to the secondary region, configure your storage account to use read-access geo-redundant storage (RA-GRS) or read-access geo-zone-redundant storage (RA-GZRS).

Box 2: No

Box 3: Yes

Box 4: Yes

Azure Cosmos DB supports multi-region writes.

NEW QUESTION 223

Select the answer that correctly completes the sentence.

Answer Area

In a relational database, each row in a table has

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer as below.

Answer Area

In a relational database, each row in a table has

NEW QUESTION 226

Your company has a reporting solution that has paginated reports. The reports query a dimensional model in a data warehouse. Which type of processing does the reporting solution use?

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

Answer: D

NEW QUESTION 227

Select the answer that correctly completes the sentence.

A

line chart

matrix

scatter

treemap

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, email Description automatically generated
A treemap chart divides the chart area into rectangles that represent the different levels and relative sizes of the data hierarchy.

NEW QUESTION 230

Match the Azure services to appropriate requirements.
To answer, drag the appropriate services from the column on the left to its requirement on the right. Each service may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.

Azure Services

Azure Data Factory

Azure Data Lake Storage

Azure SQL Database

Azure Synapse Analytics

Answer Area

Azure Service

Output data to Parquet format.

Azure Service

Store data that is in Parquet format.

Azure Service

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Azure 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.

NEW QUESTION 231

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

What should you use to query the files?

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

Answer: C

NEW QUESTION 235

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.

Customer is [answer choice].

a nested array

a nested object

a root object

Address is [answer choice].

a nested array

a nested object

a root object

Social media is [answer choice].

a nested array

a nested object

a root object

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Customer is [answer choice].

a nested array

a nested object

a root object

Address is [answer choice].

a nested array

a nested object

a root object

Social media is [answer choice].

a nested array

a nested object

a root object

NEW QUESTION 239

A data engineer is responsible for which task?

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

Answer: A

NEW QUESTION 242

.....

Relate Links

100% Pass Your DP-900 Exam with Examible Prep Materials

<https://www.exambible.com/DP-900-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>