

Microsoft

Exam Questions DP-900

Microsoft Azure Data Fundamentals



NEW QUESTION 1

Select the answer that correctly completes the sentence.

Select the answer that correctly completes the sentence.

Answer Area

Microsoft Power BI Desktop

Microsoft Power BI Desktop

The Microsoft Power BI service

The Microsoft Power BI phone app

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text, letter Description automatically generated

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

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

Creating closed caption text for audio files

is an example of

cognitive

descriptive

predictive

prescriptive

analytics.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated with low confidence

Descriptive, to answer the question: What's happening? Company name Description automatically generated



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

NEW QUESTION 5

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

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

	▼
cognitive	
descriptive	
predictive	
prescriptive	

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Text, table Description automatically generated

Generally speaking, data analytics comes in four types:

- Descriptive, to answer the question: What's happening?
- Diagnostic, to answer the question: Why's happening?
- Predictive, to answer the question: What will happen?
- Prescriptive, to answer the question: What actions should we take?



NEW QUESTION 6

You have the following SQL query.

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

NOTE: Each correct selection is worth one point.

Answer Area

Dbg.Products: A column
A database
A table
An index

ProductName: A column
A database
A table
An index

- ### A. Mastered

B. Not Mastered

Answer: A

Explanation:
Answer Area

Question 6: Which of the following is not a valid SQL identifier?

Options:

- A. A column
- B. A database
- C. A table
- D. An index

Correct Answer: B. A database

NEW QUESTION 7

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
Answer Area

Question 7: In batch processing, _____ delay in delivering data processing results is acceptable.

Options:

- A. A small
- B. A large
- C. A moderate
- D. A significant

Correct Answer: A. A small

NEW QUESTION 8

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

Statements	Yes	No
Azure Databricks is an Apache Spark-based collaborative analytics platform.	<input type="radio"/>	<input type="radio"/>
Azure Analysis Services is used for transactional workloads.	<input type="radio"/>	<input type="radio"/>
Azure Data Factory orchestrates data integration workflows.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, email Description automatically generated

NEW QUESTION 9

Select the answer that correctly completes the sentence.

Question 9: _____ natively support the analysis of relationships between entities.

Options:

- A. Column family databases
- B. Document databases
- C. Graph databases
- D. Key-value stores

Correct Answer: C. Graph databases

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Column family databases

Column family databases

Document databases

Graph databases

Key-value stores

natively support the analysis of relationships between entities.

NEW QUESTION 10

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 10

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

cognitive

descriptive

predictive

prescriptive

analytics.

- 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 11

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

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

Answer: A

Explanation:

Data Manipulation Language (DML) statements:

- > DELETE
- > INSERT
- > UPDATE

NEW QUESTION 12

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 15

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 18

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: 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 25

Match the types of data stores to the appropriate scenarios.

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

NOTE: Each correct match is worth one point.

Data Store Types	Answer Area
Graph	<div></div> Application users and their default language
Object	<div></div> Medical images and their associated metadata
Key/value	<div></div> Employee data that shows the relationships between employees

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Data Store Types	Answer Area
Graph	Key/value Application users and their default language
Object	Object Medical images and their associated metadata
Key/value	Graph Employee data that shows the relationships between employees

NEW QUESTION 26

Your company is designing an application that will write a high volume of JSON data and will have an application-defined schema. Which type of data store should you use?

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

Answer: A

Explanation:

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

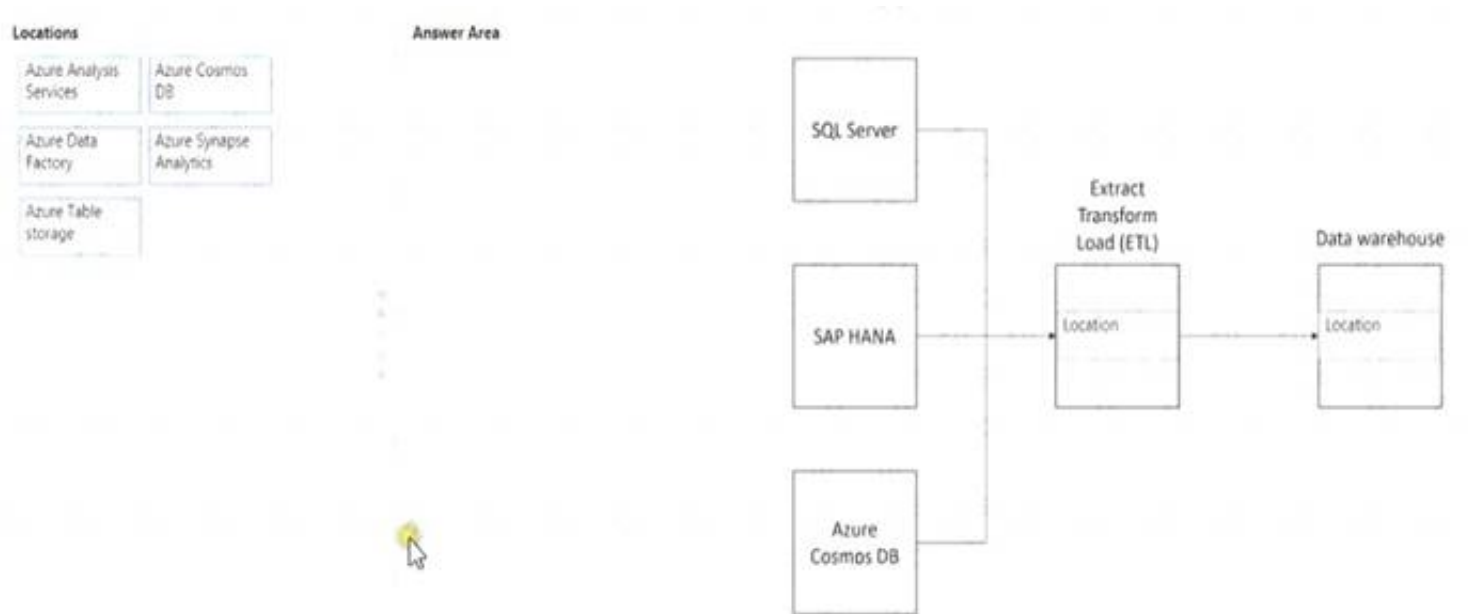
An application can store arbitrary data as a set of values. Any schema information must be provided by the application. The key/value store simply retrieves or stores the value by key.

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.

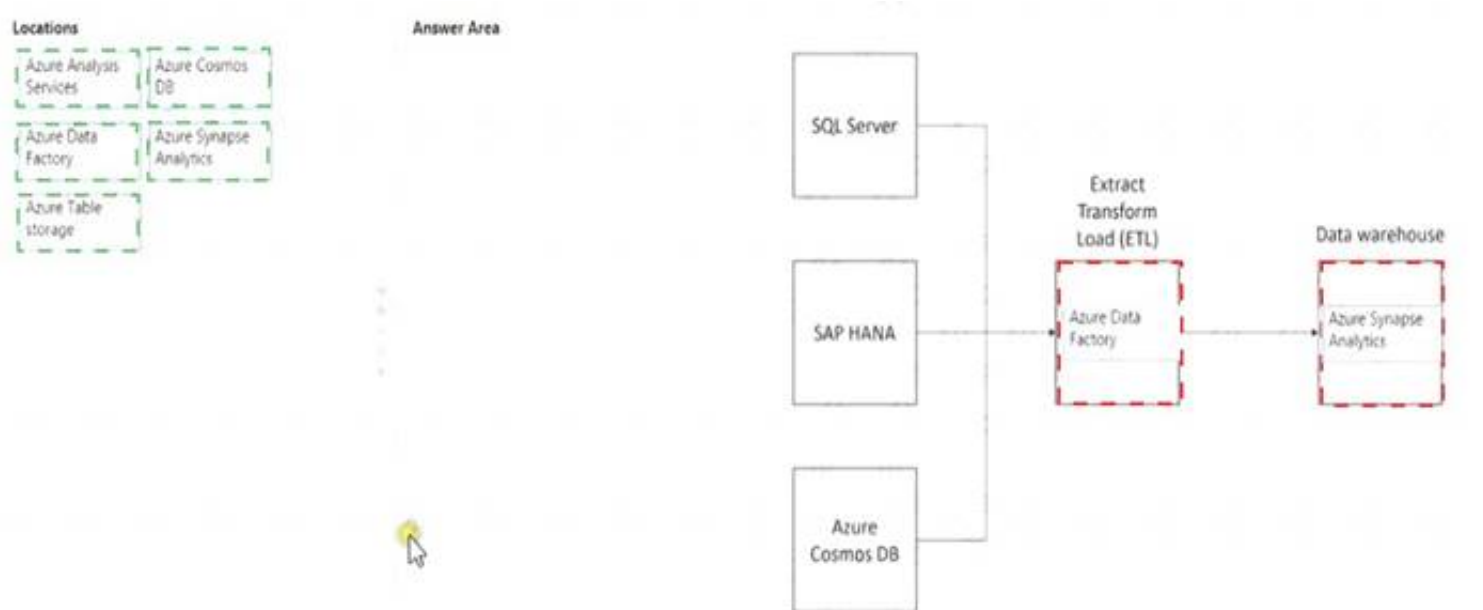
NOTE: Each correct match is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 29

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

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

Answer: C

Explanation:

Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network. Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary.

NEW QUESTION 31

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A non-relational database stores data without a fixed schema

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.

APIs

Answer Area

Cassandra API

Gremlin API

MongoDB API

Table API

Graph data

JSON documents

Key/value data

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

APIs

Answer Area

Cassandra API

Gremlin API

MongoDB API

Table API

Graph data

JSON documents

Key/value data

NEW QUESTION 37

Select the answer that correctly completes the sentence.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 39

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

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

	▼
graph.	
table.	
partition key.	
document.	

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

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

	▼
graph.	
table.	
partition key.	
document.	

NEW QUESTION 41

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

NOTE: Each correct selection is worth one point.

Statements

Yes

No

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

☐☐

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

☐☐

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

☐☐

- A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 43

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

▼

presents content defined by a query.

A heap

A stored procedure

A view

An index

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
A View

NEW QUESTION 47
Select the answer that correctly completes the sentence.

Answer Area

Select a feature mode

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Select a feature mode

When provisioning an Azure Cosmos DB

account

▼ , you need to specify which type of API you will use.

NEW QUESTION 52
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 55
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated

NEW QUESTION 60
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 64
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 66
What Is a benefit of hosting a database on Azure SQL managed instance as compared to an Azure SQL database?

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

Answer: A

NEW QUESTION 67
To complete the sentence, select the appropriate option in the answer area.

Answer Area

A key/value data store is optimized for

enforcing constraints.

simple lookups.

table joins.

transactions.

- A. Mastered

B. Not Mastered

Answer: A

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.

NEW QUESTION 71

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 76

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 80

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

Azure SQL services

Azure SQL Database elastic pool

Azure SQL Database single database

Azure SQL Managed Instance

Answer Area

A serverless configuration

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

A low-cost, low-maintenance database

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Azure SQL services

Azure SQL Database elastic pool	
Azure SQL Database single database	
Azure SQL Managed Instance	

Answer Area

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

NEW QUESTION 82

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 86

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

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

Answer: D

NEW QUESTION 87

A bank has a system that manages financial transactions.

When transferring money between accounts, the system must never retrieve a value for the source account that reflects the balance before the transfer and a value for the destination account that reflects the balance after the transfer.

Of which ACID semantic is this an example?

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

Answer: B

Explanation:

Isolation. See course material:

o concurrent transactions cannot interfere with one another, and must result in a consistent database state. o For example, while the transaction to transfer funds from one account to another is in-process, another transaction that checks the balance of these accounts must return consistent results - the balance-checking transaction can't retrieve a value for one account that reflects the balance before the transfer, and a value for the other account that reflects the balance after the transfer.

NEW QUESTION 92

Your company plans to load data from a customer relationship management (CRM) system to a data warehouse by using an extract load, and transform (ELT) process.

Where does data processing occur for each stage of the ELT process? To answer, drag the appropriate locations to the correct. Each location may be used once,

or not at all, You may need to drag the split bar between panes or scroll to view content.
 NOTE: Each correct selection is worth one point.

Locations

An in-memory data integration tool

The CRM system

The data warehouse

Answer Area

Extract: Location

Load: Location

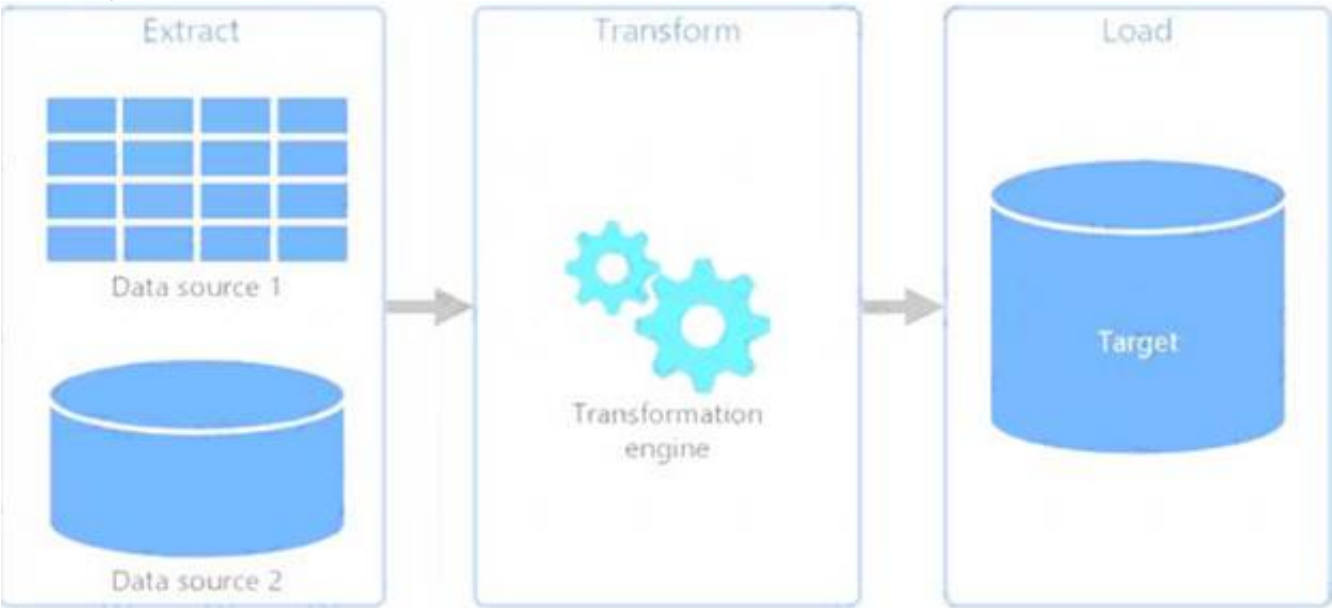
Transform: Location

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: The CRM system
 Data is extracted from the CRM system. Box 2: The data warehouse
 Data is loaded to the data warehouse.
 Box 3: An in-memory data integration tool
 The data transformation that takes place usually involves various operations, such as filtering, sorting, aggregating, joining data, cleaning data, deduplicating, and validating data.



NEW QUESTION 95

You plan to deploy an app. The app requires a nonrelational data service that will provide latency guarantees of less than 10-ms for reads and writes. What should you include in the solution?

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

Answer: A

NEW QUESTION 97

Select the answer that correctly completes the sentence.

Answer Area

A data engineer

A data analyst

A data engineer

A data scientist

A database administrator

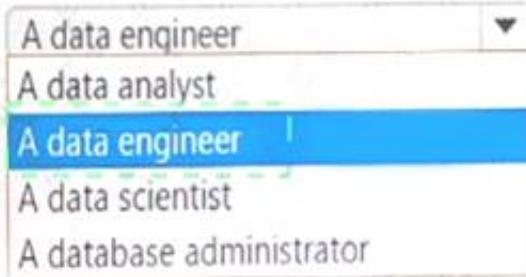
is responsible for creating visuals and charts that help a company make informed decisions.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



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

NEW QUESTION 101

Which two activities can be performed entirely by using the Microsoft Power BI service? Each correct answer presents a complete solution. (Choose two.)
NOTE: Each correct selection is worth one point.

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

Answer: AD

NEW QUESTION 102

You need to recommend a data store service that meets the following requirements:

- Native SQL API access
- Configurable indexes What should you recommend?

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

Answer: D

Explanation:

Azure Cosmos DB comes with native Core (SQL) API support.

In Azure Cosmos DB, data is indexed following indexing policies that are defined for each container. The default indexing policy for newly created containers enforces range indexes for any string or number. This policy can be overridden with your own custom indexing policy.

NEW QUESTION 106

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 111

You are deploying a software as a service (SaaS) application that requires a relational database for Online Transaction Processing (OLTP). Which Azure service should you use to support the application?

- A. Azure Cosmos D3
- B. Azure Synapse Analytics
- C. Azure HCMmtghi
- D. Azure SQL Database


Answer: D

Explanation:

Azure SQL Database is relational database and a managed service.

NEW QUESTION 115

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



Which type of analysis is shown for the fourth quarter?

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

Answer: A

NEW QUESTION 118

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 119

You have data saved in the following format:

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

Which format was used?

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

Answer: B

NEW QUESTION 123

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

NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

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.

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

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

Explanation:

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

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

Select the answer that correctly completes the sentence.

- A. Mastered
B. Not Mastered

Explanation:

Answer Area

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

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 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.

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 133
For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Statements	Yes	No
A Microsoft Power BI dashboard is associated with a single workspace.	<input type="radio"/>	<input type="radio"/>
A Microsoft Power BI dashboard can only display visualizations from a single dataset.	<input type="radio"/>	<input type="radio"/>
A Microsoft Power BI dashboard can display visualizations from a Microsoft Excel workbook.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
A Microsoft Power BI dashboard is associated with a single workspace.	<input type="radio"/>	<input checked="" type="radio"/>
A Microsoft Power BI dashboard can only display visualizations from a single dataset.	<input type="radio"/>	<input checked="" type="radio"/>
A Microsoft Power BI dashboard can display visualizations from a Microsoft Excel workbook.	<input checked="" type="radio"/>	<input type="radio"/>

NEW QUESTION 138

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A relational database must be used when

strong consistency guarantees are required.

NEW QUESTION 139

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

Data Types

Image files

Key/value pairs

Relationships between employees

Answer Area

Data Type

Azure Blob storage

Data Type

Azure Cosmos DB Gremlin API

Data Type

Azure Table storage

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Image files
 Azure Blob storage is suitable for image files. Box 2:Key/value pairs
 Azure CosmosDB table API is a key-value storage hosted in the cloud. Box 3: Relationship between employees
 One-to-many relationships between business domain objects occur frequently: for example, one department has many employees. There are several ways to implement one-to-many relationships in the Azure Table service.

NEW QUESTION 141

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

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

Answer: A

Explanation:

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

NEW QUESTION 143

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 147

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 150

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

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

Answer: D

NEW QUESTION 151

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

Answer Area

Statements	Yes	No
The Azure Cosmos DB API is configured separately for each database in an Azure Cosmos DB account.	<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"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 156

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

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

Answer: A

NEW QUESTION 160

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

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

Answer: A

NEW QUESTION 162

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 163

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

A screenshot of a computer Description automatically generated with medium confidence

NEW QUESTION 167

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

A relational database must be used when

a dynamic schema is required.

data will be stored as key/value pairs.

storing large images and videos.

strong consistency guarantees are required.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

A relational database must be used when

a dynamic schema is required.

data will be stored as key/value pairs.

storing large images and videos.

strong consistency guarantees are required.

NEW QUESTION 168

Select the answer that correctly completes the sentence.

Answer Area

Microsoft SQL Server database tables store

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Microsoft SQL Server database tables store

NEW QUESTION 170

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 173

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

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

Answer: B

NEW QUESTION 178

Your company needs to implement a relational database in Azure. The solution must minimize ongoing maintenance. Which Azure service should you use?

- A. Azure HDInsight
- B. Azure SQL Database
- C. Azure Cosmos DB
- D. SQL Server on Azure virtual machines

Answer: B

NEW QUESTION 183

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 184

You need to store data by using Azure Table storage. What should you create first?

- A. an Azure Cosmos DB instance
- B. a storage account
- C. a blob container
- D. a table

Answer: B

Explanation:

First create an Azure storage account, then use Table service in the Azure portal to create a table.
Note: An Azure storage account contains all of your Azure Storage data objects: blobs, files, queues, and tables.

NEW QUESTION 186

Match the types of analytics that can be used to answer the business questions.
To answer, drag the appropriate analytics type from the column on the left to its question on the right. Each analytics type may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.

Analytics Types

Cognitive

Descriptive

Diagnostic

Predictive

Prescriptive

Answer Area

Analytics Type

Why did sales increase last month?

Analytics Type

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

Analytics Type

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Diagnostic

Why did sales increase last month?

Predictive

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

Cognitive

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

NEW QUESTION 191

You have data saved in the following format.
{

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

Which format was used?

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

Answer: D

NEW QUESTION 193

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 198

You need to create an Azure resource to store data in Azure Table storage. Which command should you run?

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

Answer: D

NEW QUESTION 199

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

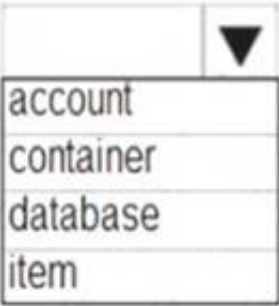
Answer: A

Explanation:

Text Description automatically generated

NEW QUESTION 203

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.


- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
Table Description automatically generated

NEW QUESTION 206
To complete the sentence, select the appropriate option in the answer area.

Answer Area

Relational data uses  to enforce relationships between different tables.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Relational data uses  to enforce relationships between different tables.

NEW QUESTION 208
You need to design and model a database by using a graphical tool that supports project-oriented offline database development. What should you use?

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

Answer: A

NEW QUESTION 209
Match the Azure services to the appropriate requirements.
To answer, drag the appropriate service from the column on the left to its requirement on the right. Each service may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.

Services	Answer Area
	 Output data to Parquet format
	 Store data that is in Parquet format
	 Persist a tabular representation of data that is stored in Parquet format
	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Azure Data Factory

Box 2: Azure Data Lake Storage

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

Box 3: Azure Synapse Analytics

Use Azure Synapse Analytics Workspaces.

NEW QUESTION 211

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

<div><div></div><div>A view</div><div>A clustered index</div><div>A stored procedure</div><div>A nonclustered index</div></div>	physically sorts the data in a table based on the values in a specified column.
---	---

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

NEW QUESTION 213

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

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

What should the developers use ?

A. Azure Data Studio

B. sqlcmd

C. Microsoft SQL Server Management Studio (SSMS)

D. Azure Data Explorer

Answer: A

Explanation:

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

Azure Data Studio offers a modern editor experience with IntelliSense, code snippets, source control integration, and an integrated terminal.

NEW QUESTION 217

You have a SQL query that combines customer data and order data. The query includes calculated columns. You need to persist the SQL query so that other users can use the query. What should you create?

A. a table

B. an index

C. scalar function

D. a view

Answer: D

Explanation:

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

NEW QUESTION 221

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 222

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

Answer Area

You can query a graph database in Azure Cosmos DB

as a JSON document by using a SQL-like language.

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:

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.

NEW QUESTION 225

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes
Azure Databricks can consume data from SQL Databases using JDBC and from SQL Databases using the Apache Spark connector. The Apache Spark connector for Azure SQL Database and SQL Server enables these databases to act as input data sources and output data sinks for Apache Spark jobs.
Box 2: Yes
You can stream data into Azure Databricks using Event Hubs. Box 3: Yes
You can run Spark jobs with data stored in Azure Cosmos DB using the Cosmos DB Spark connector. Cosmos can be used for batch and stream processing, and as a serving layer for low latency access.
You can use the connector with Azure Databricks or Azure HDInsight, which provide managed Spark clusters on Azure.

NEW QUESTION 228

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

Answer Area

Relational data is stored in

a file system as unstructured data.

a hierarchal folder structure.

a tabular form of rows and columns.

comma-separated value (CSV) files.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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.

NEW QUESTION 231

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 233

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

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

Answer: AD

NEW QUESTION 237

Select the answer that correctly completes the sentence.

Answer Area

In a relational database, the

primary key

foreign key

primary key

WHERE clause

nonclustered index

uniquely identifies each row in a table.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

In a relational database, the

primary key

foreign key

primary key

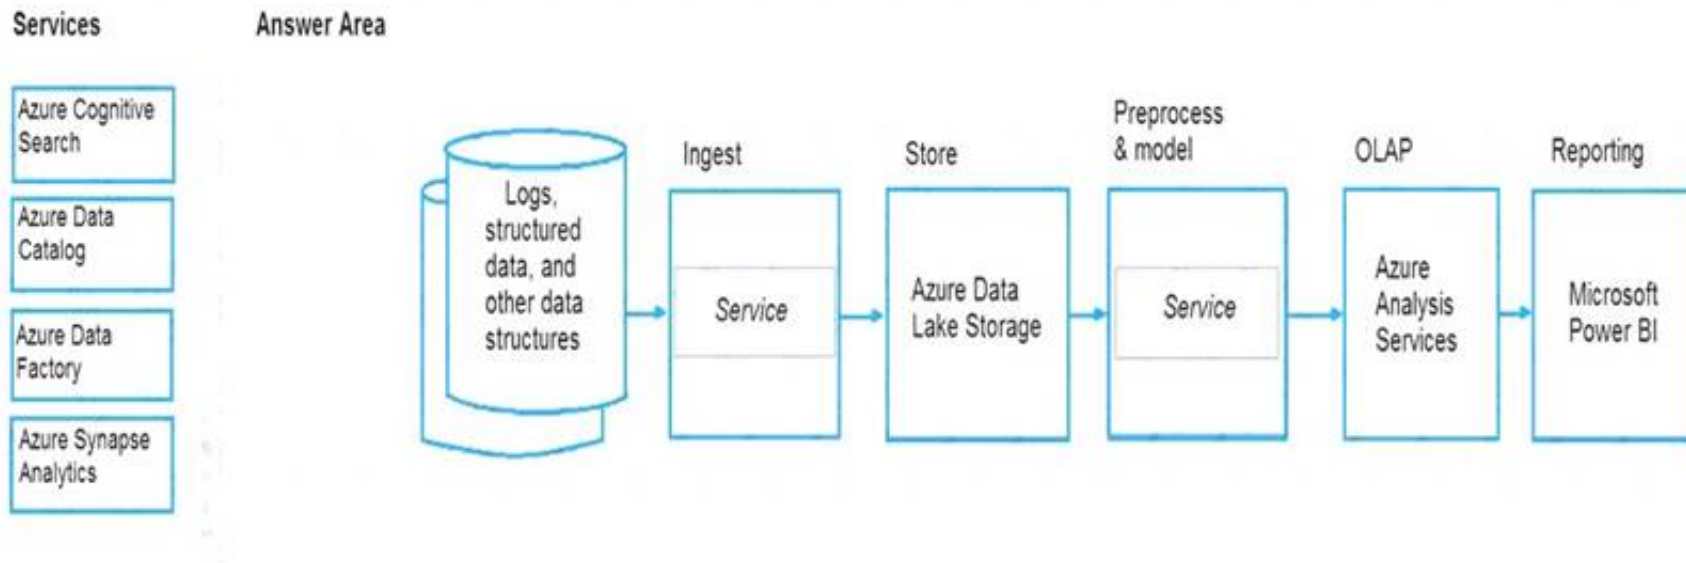
WHERE clause

nonclustered index

uniquely identifies each row in a table.

NEW QUESTION 242

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 246

What is a characteristic of batch processing?

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

Answer: D

NEW QUESTION 250

You have a quality assurance application that reads data from a data warehouse. Which type of processing does the application use?

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

Answer: A

Explanation:

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

NEW QUESTION 255

Select the answer that correctly completes the sentence.

Answer Area

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

optimized for transactions. ▼

limited in size.

optimized for transactions. |

pre-aggregated.

read-only.

NEW QUESTION 256

Select the answer that correctly completes the sentence.

Answer Area

Azure Synapse Analytics ▼

Azure SQL Database

Azure Synapse Analytics

Azure Database for MariaDB

SQL Server on Azure virtual machines

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Azure Synapse Analytics ▼

Azure SQL Database

Azure Synapse Analytics |

Azure Database for MariaDB

SQL Server on Azure virtual machines

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

NEW QUESTION 258

Select the answer that correctly completes the sentence.

Answer Area

The massively parallel processing (MPP) engine

of Azure Synapse Analytics

 ▼

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

The massively parallel processing (MPP) engine

of Azure Synapse Analytics

distributes processing across compute nodes.

 ▼

NEW QUESTION 261

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 262

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 267

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 268

At which two levels can you set the throughput for an Azure Cosmos DB account? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. container
- B. item

- C. database
- D. partition

Answer: AC

NEW QUESTION 273

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 277

Select the answer that correctly completes the sentence.

Answer Area

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

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

a matching schema in the data source and the data target.

a target data store powerful enough to transform data.

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

that the data target be a relational database.

Content Display

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

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

a matching schema in the data source and the data target.

a target data store powerful enough to transform data.

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

that the data target be a relational database.

Content Display

NEW QUESTION 281

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

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

Answer: BC

Explanation:

Azure SQL Database is a fully managed platform as a service (PaaS) database engine that handles most of the database management functions such as upgrading, patching, backups, and monitoring without user involvement. SQL Database delivers predictable performance with multiple resource types, service tiers, and compute sizes. It provides dynamic scalability with no downtime,

built-in intelligent optimization, global scalability and availability, and advanced security options. These capabilities allow you to focus on rapid app development and accelerating your time-to-market, rather than on managing virtual machines and infrastructure.

NEW QUESTION 284

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 285

Select the answer the correctly completes the sentence.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A database administrator

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

NEW QUESTION 289

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

In batch processing,

data is always inserted one row at a time.

data is processed in real-time.

latency is expected.

processing car only execute serially.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

In batch processing,

data is always inserted one row at a time.

data is processed in real-time.

latency is expected.

processing car only execute serially.

NEW QUESTION 290

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 292

Select the answer that correctly completes the sentence.

Answer Area

A block of code that runs in a database is called

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A block of code that runs in a database is called

a stored procedure.

NEW QUESTION 295

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	
			<div>YesNo</div>
		Batch processing can output data to a file store.	<div><div></div><div></div></div>
		Batch processing can output data to a relational database.	<div><div></div><div></div></div>
		Batch processing can output data to a NoSQL database.	<div><div></div><div></div></div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area			
		Statements	
			<div>YesNo</div>
		Batch processing can output data to a file store.	<div><div><div></div></div><div></div></div>
		Batch processing can output data to a relational database.	<div><div><div></div></div><div></div></div>
		Batch processing can output data to a NoSQL database.	<div><div></div><div><div></div></div></div>

NEW QUESTION 298

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

Answer Area	
An extract, load, and transform (ELT) process requires	<div><div>a data pipeline that includes a transformation engine.</div><div>a separate transformation engine.</div><div>a target data store powerful enough to transform data.</div><div>data that is fully processed before being loaded to the target data store.</div></div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area	
An extract, load, and transform (ELT) process requires	<div><div>a data pipeline that includes a transformation engine.</div><div>a separate transformation engine.</div><div>a target data store powerful enough to transform data.</div><div>data that is fully processed before being loaded to the target data store.</div></div>

NEW QUESTION 300

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

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

Answer: D

Explanation:

Time series data is a set of values organized by time. Time series databases typically collect large amounts of data in real time from a large number of sources. Updates are rare, and deletes are often done as bulk operations. Although the records written to a time-series database are generally small, there are often a large number of records, and total data size can grow rapidly.

NEW QUESTION 303

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 306

You manage an application that stores data in a shared folder on a Windows server. You need to move the shared folder to Azure Storage. Which type of Azure Storage should you use?

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

Answer: C

Explanation:

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

NEW QUESTION 308

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 309

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

DP-900 Practice Exam Features:

- * DP-900 Questions and Answers Updated Frequently
- * DP-900 Practice Questions Verified by Expert Senior Certified Staff
- * DP-900 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * DP-900 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The DP-900 Practice Test Here](#)