

➤ **Vendor: Microsoft**

➤ **Exam Code: DP-900**

➤ **Exam Name: Microsoft Azure Data Fundamentals**

➤ **New Updated Questions from [Braindump2go](#) (Updated in [February/2021](#))**

[Visit Braindump2go and Download Full Version DP-900 Exam Dumps](#)

QUESTION 111

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: C

QUESTION 112

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

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

Answer: A

Explanation:

Data Manipulation Language (DML) statements:

DELETE

INSERT

UPDATE

Reference:

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

QUESTION 113

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: B

QUESTION 114

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

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

Answer: B

QUESTION 115

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

QUESTION 116

Your company needs to ensure that an 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

[DP-900 Exam Dumps](#) [DP-900 Exam Questions](#) [DP-900 PDF Dumps](#) [DP-900 VCE Dumps](#)

<https://www.braindump2go.com/dp-900.html>

D. Azure Traffic Manager

Answer: B

QUESTION 117

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

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

Answer: B

QUESTION 118

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

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

Answer: A

Explanation:

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

Reference:

<https://www.sqlshack.com/configuring-the-azure-sql-database-firewall/>

QUESTION 119

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

- A. the Azure Active directory (Azure AD) account that created the database
- B. the Azure Active Directory (Azure AD) administrator account
- C. the sa account
- D. the server admin login account of the logical server

Answer: B

QUESTION 120

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: B

QUESTION 121

Hotspot Question

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

Answer Area

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

<input type="text"/>
changes to relationships between entities
geographically distributed writes
transactional writes
writes that have varying data structures

Answer:

Answer Area

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

<input type="text"/>
changes to relationships between entities
geographically distributed writes
transactional writes
writes that have varying data structures

Explanation:

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

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

Incorrect Answers:

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

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

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

QUESTION 122

Hotspot Question

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

Answer:

Answer Area

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

cognitive
descriptive
predictive
prescriptive

Explanation:

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

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

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

Predictive, to answer the question: What will happen?

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

Reference:

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

QUESTION 123

Drag and Drop Question

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

Where does data processing occur for each stage of the ELT process? To answer, drag the appropriate locations to the correct stages. Each location may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Locations

- An in-memory data integration tool
- The CRM system
- The data warehouse

Answer Area

Extract:

Load:

Transform:

Answer:

Locations

- An in-memory data integration tool
- The CRM system
- The data warehouse

Answer Area

Extract:

Load:

Transform:

Explanation:

Box 1: The CRM system

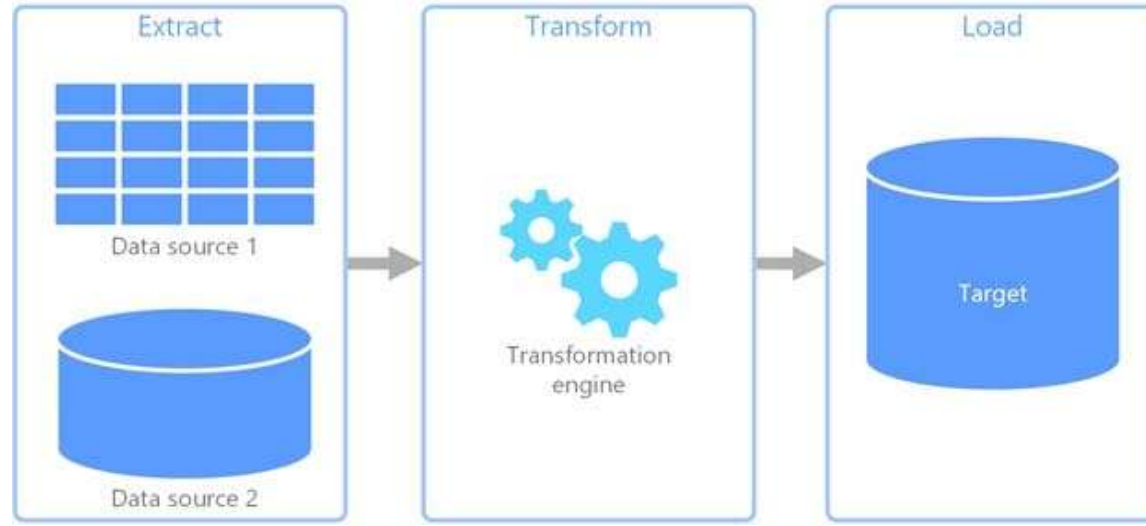
Data is extracted from the CRM system.

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.



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

QUESTION 124

Drag and Drop Question

Match the types of workloads the appropriate scenario.

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

NOTE: Each correct match is worth one point.

Workload Types	Answer Area
<div style="display: flex; justify-content: space-around;"> <div style="border: 1px solid gray; padding: 2px;">Batch</div> <div style="border: 1px solid gray; padding: 2px;">Micro-batch</div> <div style="border: 1px solid gray; padding: 2px;">Streaming</div> </div>	<div style="border: 1px solid gray; padding: 5px;"> <p>Workload Type: Data for a product catalog will be loaded every 12 hours to a data warehouse.</p> <p>Workload Type: Data for online purchases will be loaded to a data warehouse as the purchases occur.</p> <p>Workload Type: Updates to inventory data will be loaded to a data warehouse every 1,000 transactions.</p> </div>
<div style="border: 1px solid gray; padding: 2px;">Batch</div> <div style="border: 1px solid gray; padding: 2px;">Streaming</div> <div style="border: 1px solid gray; padding: 2px;">Micro-batch</div>	<div style="border: 1px solid gray; padding: 5px;"> <p>Batch: Data for a product catalog will be loaded every 12 hours to a data warehouse.</p> <p>Streaming: Data for online purchases will be loaded to a data warehouse as the purchases occur.</p> <p>Micro-batch: Updates to inventory data will be loaded to a data warehouse every 1,000 transactions.</p> </div>

Answer:

QUESTION 125

Hotspot Question

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

Note: Each correct selection is worth one point.

Answer Area

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

Answer:

Answer Area

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

Explanation:

Box 1: Yes

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

Box 2: No

Box 3: No

Reference:

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