



Vendor: Microsoft

Exam Code: DP-900

Exam Name: Microsoft Azure Data Fundamentals

Version: DEMO

QUESTION 1

Which database transaction property ensures that individual transactions are executed only once and either succeed in their entirety or roll back?

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

Answer: C

Explanation:

An atomic transaction is an indivisible and irreducible series of database operations such that either all occurs, or nothing occurs. A guarantee of atomicity prevents updates to the database occurring only partially, which can cause greater problems than rejecting the whole series outright. As a consequence, the transaction cannot be observed to be in progress by another database client.

QUESTION 2

Which property of a transactional workload guarantees that each transaction is treated as a single unit that either succeeds completely or fails completely?

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

Answer: D

Explanation:

A.C.I.D. properties: Atomicity, Consistency, Isolation, and Durability

Atomicity - each statement in a transaction (to read, write, update or delete data) is treated as a single unit. Either the entire statement is executed, or none of it is executed. This property prevents data loss and corruption from occurring if, for example, if your streaming data source fails mid-stream.

<https://www.databricks.com/glossary/acid-transactions>

QUESTION 3

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

- A. GRANT
- B. INSERT
- C. REVOKE
- D. DROP

Answer: B

Explanation:

Data Manipulation Language (DML) affect the information stored in the database. Use these statements to insert, update, and change the rows in the database.

- BULK INSERT
- DELETE
- INSERT
- SELECT

- UPDATE
- MERGE

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/statements>

QUESTION 4

Hotspot Question

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

<input type="text"/>	_____ natively support the analysis of relationships between entities.
Column family databases	
Document databases	
Graph databases	
Key-value stores	

Answer:

<input type="text"/>	_____ natively support the analysis of relationships between entities.
Column family databases	
Document databases	
Graph databases	
Key-value stores	

Explanation:

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

Reference:

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

QUESTION 5

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: A

Explanation:

A data warehouse is relational in nature. This means that the structure or schema of the data is determined by predefined business and product requirements that are curated, conformed, and optimized for SQL query operations. As a result, data warehouses are best used for storing data

that has been treated with a specific purpose in mind, such as data mining for BI analysis, or for sourcing a business use case that has already been identified.

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

QUESTION 6

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.

Reference:

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

QUESTION 7

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.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-manage-indexing-policy>

QUESTION 8

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 9

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
- 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.

Reference:

<https://techcommunity.microsoft.com/t5/azure-database-support-blog/azure-sql-db-private-link-private-endpoint-connectivity/ba-p/1235573>

QUESTION 10

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

- A. relational
- B. 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.

Reference:

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

QUESTION 11

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).

Reference:

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

QUESTION 12

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.

Reference:

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

QUESTION 13

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

Explanation:

Predictive, to answer the question: What will happen?



Reference:

<https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitive-analytics/>
<https://azure.microsoft.com/en-us/blog/answering-whats-happening-whys-happening-and-what-will-happen-with-iot-analytics/>

QUESTION 14

Drag and Drop Question

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

Answer:

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

Explanation:

Box 1: Azure SQL Database single database

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

Scenarios well suited for serverless compute

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

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

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

Box 2: Azure SQL Managed Instance

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

Box 3: Azure SQL Database elastic pool

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

Reference:

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

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

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

QUESTION 15

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

NOTE: Each correct selection is worth one point.

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

Answer: BD

Explanation:

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

Reference:

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

Thank You for Trying Our Product

Braindump2go Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.braindump2go.com/all-products.html>



10% Discount Coupon Code: ASTR14