

➤ **Vendor: Microsoft**

➤ **Exam Code: DP-900**

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

➤ **New Updated Questions from [Braindump2go](#) (Updated in [June/2023](#))**

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

**QUESTION 104**

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 105**

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.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/overview>

**QUESTION 106**

When provisioning an Azure Cosmos DB account, which feature provides redundancy within an Azure region?

- A. multi-master replication
- B. Availability Zones
- C. automatic failover
- D. the strong consistency level

**Answer: B**

**Explanation:**

With Availability Zone (AZ) support, Azure Cosmos DB will ensure replicas are placed across multiple zones within a given region to provide high availability and resiliency to zonal failures.

Note: Azure Cosmos DB provides high availability in two primary ways. First, Azure Cosmos DB replicates data across regions configured within a Cosmos account. Second, Azure Cosmos DB maintains 4 replicas of data within a region.

Reference:

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

**QUESTION 107**

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**

**QUESTION 108**

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

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

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: B**

**Explanation:**

<https://learn.microsoft.com/en-us/azure/architecture/guide/technology-choices/images/document.png>

**QUESTION 109**

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 110**

What is a benefit of the Azure Cosmos DB Table API as compared to Azure Table storage?

- A. supports partitioning
- B. provides resiliency if an Azure region fails
- C. provides a higher storage capacity
- D. supports a multi-master model

**Answer: D**

**Explanation:**

Multi-master support for Azure Cosmos DB is now available in all public regions.

Azure CosmosDB table API is a key-value storage hosted in the cloud. It's a part of Azure Cosmos DB, that is Microsoft's multi-model database.

Reference:

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

**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-production account type
- D. API

**Answer: D**

**Explanation:**

You must create a separate account for each API.

Reference:

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

**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) affects 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 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?

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

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

- 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.

Reference:

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

**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: BD**

**Explanation:**

B: 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.

D: 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.

Reference:

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

**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
- 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 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 HDInsight
- D. Azure SQL Database

**Answer: D**

**Explanation:**

Azure SQL Database is relational database and a managed service.

Incorrect Answers:

A, C: Cosmos DB, HDInsight are non-relational databases.

B: Azure Synapse Analytics is for data warehousing, not for Online Transaction Processing.

Reference:

<https://cloud.netapp.com/blog/azure-cvo-blg-azure-database-review-your-guide-for-database-assessment>

**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/config-ring-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 director/ (Azure AO) account that created the database
- B. the Azure Active Directly Azure AD administrator account
- C. the sa account
- D. the server admin login account of the logical server

**Answer: D**

**Explanation:**

When you first deploy Azure SQL, you specify an admin login and an associated password for that login. This administrative account is called Server admin.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-quickstart>

**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 each of their 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.

Reference:

<https://docs.microsoft.com/en-us/sql/azure-data-studio/download-azure-data-studio>

**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

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

**Answer:**

**Answer Area**

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

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

**Explanation:**

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

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

Incorrect Answers:

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

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

Reference:

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

**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	Answer Area
An in-memory data integration tool	Extract: Location
The CRM system	Load: Location
The data warehouse	Transform: Location

**Answer:**

Locations	Answer Area
An in-memory data integration tool	Extract: The CRM system
The CRM system	Load: The data warehouse
The data warehouse	Transform: The data warehouse

**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: The data warehouse

In the ELT pipeline, the transformation occurs in the target data store. Instead of using a separate transformation engine, the processing capabilities of the target data store are used to transform data.

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
Batch	Workload Type: Data for a product catalog will be loaded every 12 hours to a data warehouse.
Micro-batch	Workload Type: Data for online purchases will be loaded to a data warehouse as the purchases occur.
Streaming	Workload Type: Updates to inventory data will be loaded to a data warehouse every 1,000 transactions.

**Answer:**

Worksheet Types

Answer Area

Batch	Data for a product catalog will be loaded every 12 hours to a data warehouse.
Streaming	Data for online purchases will be loaded to a data warehouse as the purchases occur.
Micro-batch	Updates to inventory data will be loaded to a data warehouse every 1,000 transactions.

**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 checked="" type="radio"/>	<input type="radio"/>
Batch processing can output data to a NoSQL database	<input checked="" type="radio"/>	<input type="radio"/>

**Explanation:**

Box 1: Yes

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

Box 2: Yes

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

Box 3: Yes

Reference:

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

**QUESTION 143**

Which T-SQL statement should be used to instruct a database management system to use an index instead of performing a full table scan?

- A. SELECT
- B. WHERE
- C. JOIN

**Answer: A**

**Explanation:**

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

This applies to the following commands:

DELETE

INSERT

SELECT

UPDATE

MERGE

Reference:

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

**QUESTION 144**

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**

**QUESTION 145**

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**

**Explanation:**

<https://learn.microsoft.com/en-us/sql/relational-databases/performance/joins?view=sql-server-ver16>

**QUESTION 146**

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**

**QUESTION 147**

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).

Reference:

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

**QUESTION 151**

What do you use to create Power BI paginated reports?

- A. Power BI service
- B. Power BI Report Builder
- C. Power BI Dashboard
- D. Power BI Desktop

**Answer: B**

**Explanation:**

You use the Power BI Report Builder to author and publish paginated reports. You create a paginated report by creating a report definition that specifies what data to retrieve, where to get it, and how to display it. The report is generated by the report processor when you run the report. You can preview the report in Report Builder before publishing it to the Power BI service.

You should not use the Power BI Dashboard to create paginated reports. Power BI Dashboard is a single page on which your visualizations are posted as tiles on the dashboard to display information. A dashboard allows you to show important business metrics at a glance.

You should not use Power BI service. The Power BI service is a set of analysis and display tools that lets you create visuals based on your data. This includes Power BI Desktop and Power BI Dashboard.

You should not use Power BI Desktop, which is a component of the Power BI service. Power BI Desktop is used to create interactive reports, typically for publishing in Power BI Dashboard.

**QUESTION 152**

Which Azure storage service includes support for Hot, Cool, and Archive access tiers?

- A. Azure File
- B. Azure Blob
- C. Azure Table

**Answer: B**

**Explanation:**

Azure Blob is the only Azure storage option that supports access tiers. The default is the Hot tier, which is designed for frequently accessed data. The Cool tier is optimized for data that will be stored for at least 30 days. The Cool tier has lower storage cost than the Hot tier but higher costs for early access. The Archive tier is designed for data that is rarely accessed and will remain in storage for at least 180 days. Access to Archive tier data requires the data to be rehydrated to a Hot or Cool tier. This can mean a latency of several hours. Access tier support requires Data Lake Storage Gen2.

Azure Blob also supports two performance tiers. The Standard performance tier provides for high performance using hard disk-based storage media. The Premium performance provides greater throughput than the Standard tier and uses solid state drive (SSD) media. The Standard and Premium tiers are also supported for other storage options including Azure File storage and Azure SQL Database.

Azure Table and Azure File do not support access tiers. Access tiers is a feature supported through Cosmos DB. Azure Table and Azure File are distinct storage types and are not implemented through Cosmos DB APIs. Table storage is used for storing structured, non-relational data. File storage provides file storage with shared access, similar to a file server.

**QUESTION 153**

In which two situations should you use a data warehousing solution? Each correct answer presents a complete solution.

- A. You want to provide a platform for data mining.
- B. You want to consolidate data from multiple sources into a non-relational store.
- C. You want to optimize performance for a widely distributed transactional processing application.

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

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

D. You want to generate reports from historical data without impacting transactional processing.

**Answer:** AD

**Explanation:**

Two common use cases for data warehousing are:

- You want to generate reports from historical data without impacting transactional processing.
- You want to provide a platform for data mining.

Data warehousing lets you consolidate data from multiple sources for analysis and reporting. Data stores are optimized for read operations with few, if any, writes performed on the data. There are typically no locking requirements in a data warehouse.

You are not consolidating data into a non-relational store. Azure data warehousing solutions use relational data with a schema strongly enforced on data load and write. Data is sometimes less normalized in a data warehousing solution than data for transaction processing. The complex queries and analysis in data warehousing is often better suited to less normalized data.

You would not use a data warehousing solution for a widely distributed transactional processing application. This requires a data solution optimized for write operations that provides for data locking and strong consistency.

#### **QUESTION 154**

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

**Explanation:**

<https://docs.microsoft.com/en-us/azure/postgresql/overview-postgres-choose-server-options>

#### **QUESTION 162**

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

- A. REVOKE
- B. DISABLE
- C. CREATE
- D. UPDATE

**Answer:** D

**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 163**

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

- A. az storage share create
- B. az storage account create
- C. az cosmosdb create
- D. az storage container create

**Answer:** B

**Explanation:**

<https://docs.microsoft.com/en-us/cli/azure/storage/container?view=azure-cli-latest>

#### **QUESTION 164**

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

**Explanation:**

To add a column in a view Alter.

<https://docs.microsoft.com/en-us/sql/relational-databases/tables/add-columns-to-a-table-database-engine?view=sql-server-ver15>

Insert adds one or more rows to a table.

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

#### **QUESTION 165**

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

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



**Answer: D****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

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/statements>**QUESTION 166**

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.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/tables/table-storage-quickstart-portal><https://docs.microsoft.com/en-us/azure/storage/common/storage-account-create>**QUESTION 167**

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 168**

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 169**

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/>

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

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