



**Vendor: Microsoft**

**Exam Code: DP-900**

**Exam Name: Microsoft Azure Data Fundamentals**

**Version: 13.01**

**Q & As: 198**

**QUESTION 1**

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

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

**Correct Answer:** C

**QUESTION 2**

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

**Correct Answer:** AD

**QUESTION 3**

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

**Correct 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.

**QUESTION 4**

**DRAG DROP**

Match the Azure data factory components to the appropriate descriptions.

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

NOTE: Each correct match is worth one point.

Components	Answer Area
Dataset	Component A representation of data structures within data stores
Linked service	Component The information used to connect to external resources
Mapping data flow	Component A logical grouping of activities that performs a unit of work and can be scheduled
Pipeline	

**Correct Answer:**

Components	Answer Area
Dataset	Dataset A representation of data structures within data stores
Linked service	Linked service The information used to connect to external resources
Mapping data flow	
Pipeline	Pipeline A logical grouping of activities that performs a unit of work and can be scheduled

**QUESTION 5**

**HOTSPOT**

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 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"/>

**Correct Answer:**

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

**QUESTION 6**

**HOTSPOT**

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

An extract, transform, and load (ETL) process

requires

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.

**Correct Answer:**

An extract, transform, and load (ETL) process

requires

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.

**QUESTION 7**

**HOTSPOT**

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

**Correct Answer:**

natively support the analysis of relationships between entities.

Column family databases
Document databases
Graph databases
Key-value stores

**QUESTION 8**

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

**Correct Answer: D**

**Explanation:**

Azure SQL Database is relational database and a managed service.