DP-900^{Q&As}

Microsoft Azure Data Fundamentals

Pass Microsoft DP-900 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.leads4pass.com/dp-900.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

Instant Download After Purchase

100% Money Back Guarantee

365 Days Free Update

800,000+ Satisfied Customers

QUESTION 1

You have a transactional application that stem data in an Azure SQL managed instance. When should you implement a read-only database replica?

A. You need to generate reports without affecting the transactional workload.

B. You need to audit the transactional application.

C. You need to implement high availability in the event of a regional outage

D. You need to improve the recovery point objective (RPO).

Correct Answer: A

Use read-only replicas to offload read-only query workloads.

Reference: https://docs.microsoft.com/en-us/azure/azure-sql/database/read-scale-out

QUESTION 2

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

Hot 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 (COL)
	as a partitioned row store by using Language-integrated Query (UNO)
	as nodes and edges by using the Gremlin language

Correct Answer:

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 (COL)
	as a partitioned row store by using Language-integrated Query (UNO)
	as nodes and edges by using the Gremlin language

Reference: https://www.sqlshack.com/graph-database-implementation-with-azure-cosmos-db-using-the-api/

QUESTION 3

HOTSPOT

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

Hot Area:

A view
A clustered index
A stored procedure
A nonclustered index

physically sorts the data in a table based on the values in a specified column.

Correct Answer:

A view
A clustered index
A stored procedure
A nonclustered index

physically sorts the data in a table based on the values in a specified column.

Reference: https://docs.microsoft.com/en-us/sql/relational-databases/indexes/clustered-and-nonclustered-indexes-described?view=sql-server-ver15

QUESTION 4

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

Hot Area:

	Yes	No
If you have a platform as a service (PaaS) database in Azure, you are responsible for applying operating system updates.	○	○
If you have a platform as a service (PaaS) database in Azure, backups are performed automatically.	○	○
If you have a platform as a service (PaaS) database in Azure, you are responsible for upgrading the database engine.	○	○

Correct Answer:

	Yes	No
If you have a platform as a service (PaaS) database in Azure, you are responsible for applying operating system updates.	○	●
If you have a platform as a service (PaaS) database in Azure, backups are performed automatically.	●	○
If you have a platform as a service (PaaS) database in Azure, you are responsible for upgrading the database engine.	○	●

Box 1: No

Microsoft handles all patching and updating of the SQL and operating system code. You don\\'t have to manage the underlying infrastructure.

Box 2: Yes

SQL Database is a fully managed service that has built-in high availability, backups, and other common maintenance operations.

Box 3: No

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

---

QUESTION 5

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

Hot Area:

An extract, load, and transform (ELT) process required	▼
	a data pipeline that includes a transformation engine
	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

Correct Answer:

An extract, load, and transform (ELT) process required	▼
	a data pipeline that includes a transformation engine
	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