DP-300 Q&As

Administering Relational Databases on Microsoft Azure

Pass Microsoft DP-300 Exam with 100% Guarantee

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

https://www.pass4itsure.com/dp-300.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

DRAG DROP

Your company analyzes images from security cameras and sends alerts to security teams that respond to unusual activity. The solution uses Azure Databricks.

You need to send Apache Spark level events, Spark Structured Streaming metrics, and application metrics to Azure Monitor.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions in the answer area and arrange them in the correct order.

Select and Place:

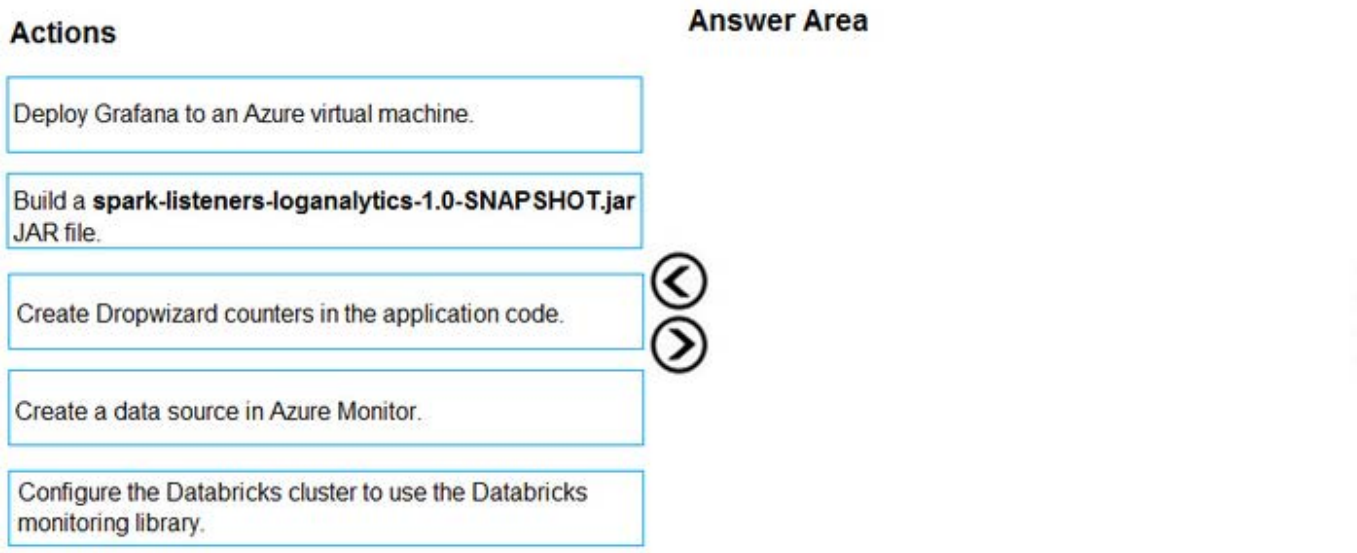

Correct Answer:

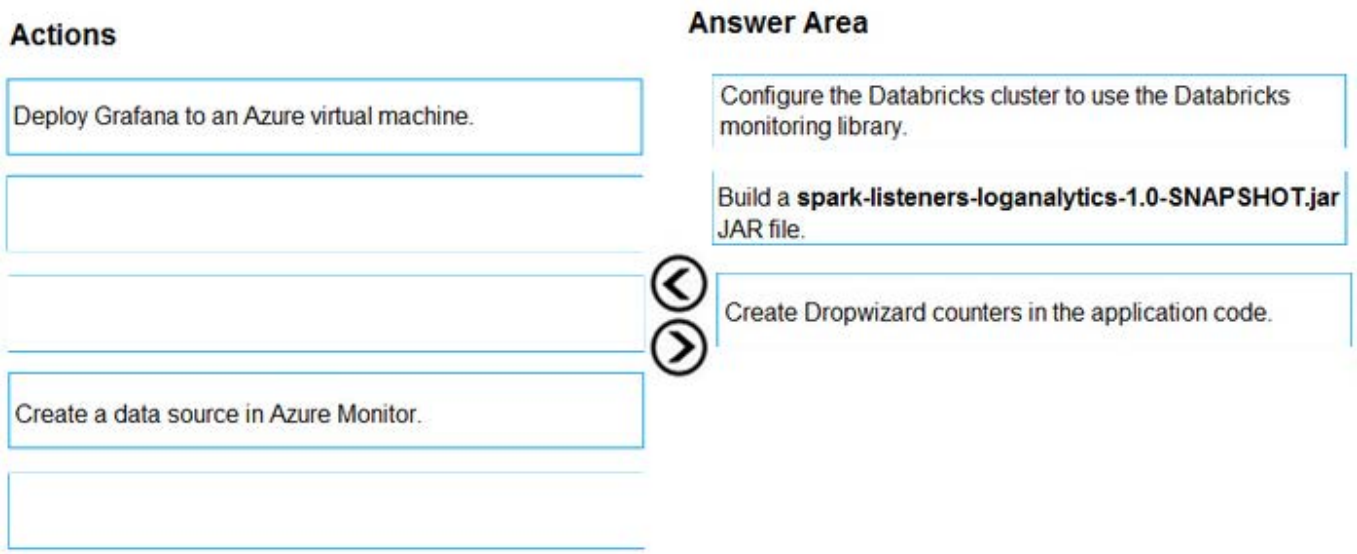

Send application metrics using Dropwizard.

Spark uses a configurable metrics system based on the Dropwizard Metrics Library.

To send application metrics from Azure Databricks application code to Azure Monitor, follow these steps:

Step 1: Configure your Azure Databricks cluster to use the Databricksmonitoring library.

Prerequisite: Configure your Azure Databricks cluster to use the monitoring library.

Step 2: Build the spark-listeners-loganalytics-1.0-SNAPSHOT.jar JAR file

Step 3: Create Dropwizard counters in your application code

Create Dropwizard gauges or counters in your application code

QUESTION 2

HOTSPOT

You have an Azure SQL database.

You are reviewing a slow performing query as shown in the following exhibit.

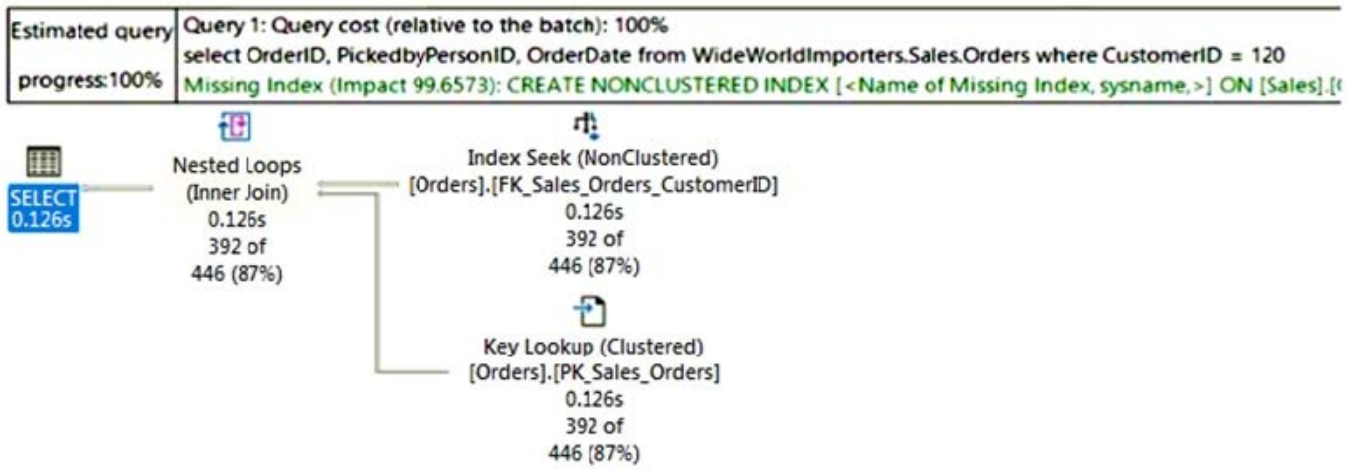

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

The exhibit shows **[answer choice]**.

- an actual execution plan
- an estimated execution plan
- Live Query Statistics

The **[answer choice]** operator in the execution plan indicates that the query would benefit from performance tuning.

- Index Seek
- Key Lookup
- Nested Loops

Correct Answer:

Answer Area

The exhibit shows **[answer choice]**.

- an actual execution plan
- an estimated execution plan
- Live Query Statistics

The **[answer choice]** operator in the execution plan indicates that the query would benefit from performance tuning.

- Index Seek
- Key Lookup
- Nested Loops

Reference: https://docs.microsoft.com/en-us/sql/relational-databases/performance/live-query-statistics?view=sql-server-ver15

QUESTION 3

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while

others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure SQL database named Sales.

You need to implement disaster recovery for Sales to meet the following requirements:

1.

During normal operations, provide at least two readable copies of Sales.

2.

Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the General Purpose service tier and geo- replication.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones.

Note: Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

Reference: https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla

QUESTION 4

You have an Azure subscription that contains an instance of SQL Server on an Azure virtual machine named SQLVM1 and a user named User1. SQLVM1 hosts a database named DB1.

You need to ensure that User1 can create a scheduled task to perform a full backup of DB1. The solution must use the principle of least privilege.

Which built-in database role should you assign to User1?

A. db_owner

B. SQLAgentReaderRole

C. SQLAgentUserRole

D. SQLAgentOperatorRole

Correct Answer: C

QUESTION 5

You need to implement the surrogate key for the retail store table. The solution must meet the sales transaction dataset requirements. What should you create?

A. a table that has a FOREIGN KEYconstraint

B. a table the has an IDENTITY property

C. a user-defined SEQUENCE object

D. a system-versioned temporal table

Correct Answer: B

Scenario: Contoso requirements for the sales transaction dataset include: Implement a surrogate key to account for changes to the retail store addresses.

A surrogate key on a table is a column with a unique identifier for each row. The key is not generated from the table data. Data modelers like to create surrogate keys on their tables when they design data warehouse models. You can use the IDENTITY property to achieve this goal simply and effectively without affecting load performance.

Reference: https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-identity