DP-200 Q&As

Implementing an Azure Data Solution

Pass Microsoft DP-200 Exam with 100% Guarantee

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

https://www.pass4itsure.com/dp-200.html

100% Passing Guarantee
100% Money Back Assurance

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

QUESTION 1

A company uses Azure Data Lake Gen 1 Storage to store big data related to consumer behavior.

You need to implement logging.

Solution: Configure an Azure Automation runbook to copy events.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

Instead configure Azure Data Lake Storage diagnostics to store logs and metrics in a storage account.

References: https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-diagnostic-logs

---

QUESTION 2

A company has a real-time data analysis solution that is hosted on Microsoft Azure. The solution uses Azure Event Hub to ingest data and an Azure Stream Analytics cloud job to analyze the data. The cloud job is configured to use 120

Streaming Units (SU).

You need to optimize performance for the Azure Stream Analytics job.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Implement event ordering

B. Scale the SU count for the job up

C. Implement Azure Stream Analytics user-defined functions (UDF)

D. Scale the SU count for the job down

E. Implement query parallelization by partitioning the data output

F. Implement query parallelization by partitioning the data input

Correct Answer: BF

Scale out the query by allowing the system to process each input partition separately.

F: A Stream Analytics job definition includes inputs, a query, and output. Inputs are where the job reads the data stream from.

References: https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-parallelization

QUESTION 3

A company runs Microsoft Dynamics CRM with Microsoft SQL Server on-premises. SQL Server Integration Services (SSIS) packages extract data from Dynamics CRM APIs, and load the data into a SQL Server data warehouse.

The datacenter is running out of capacity. Because of the network configuration, you must extract on premises data to the cloud over https. You cannot open any additional ports. The solution must implement the least amount of effort.

You need to create the pipeline system.

Which component should you use? To answer, select the appropriate technology in the dialog box in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

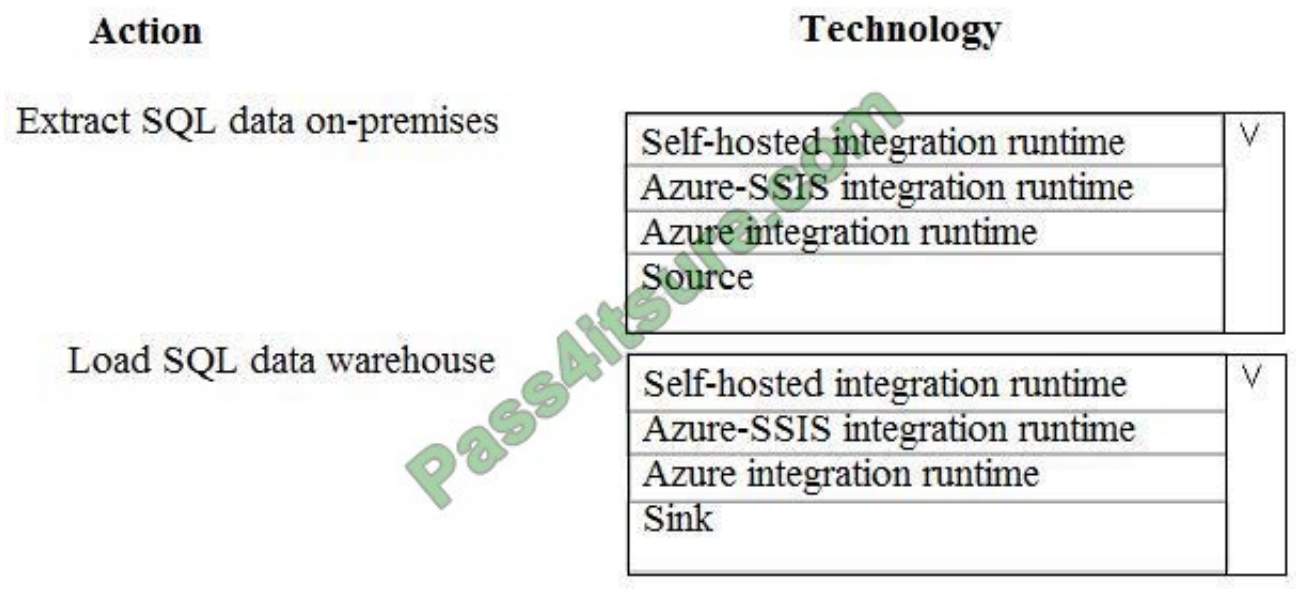

Correct Answer:

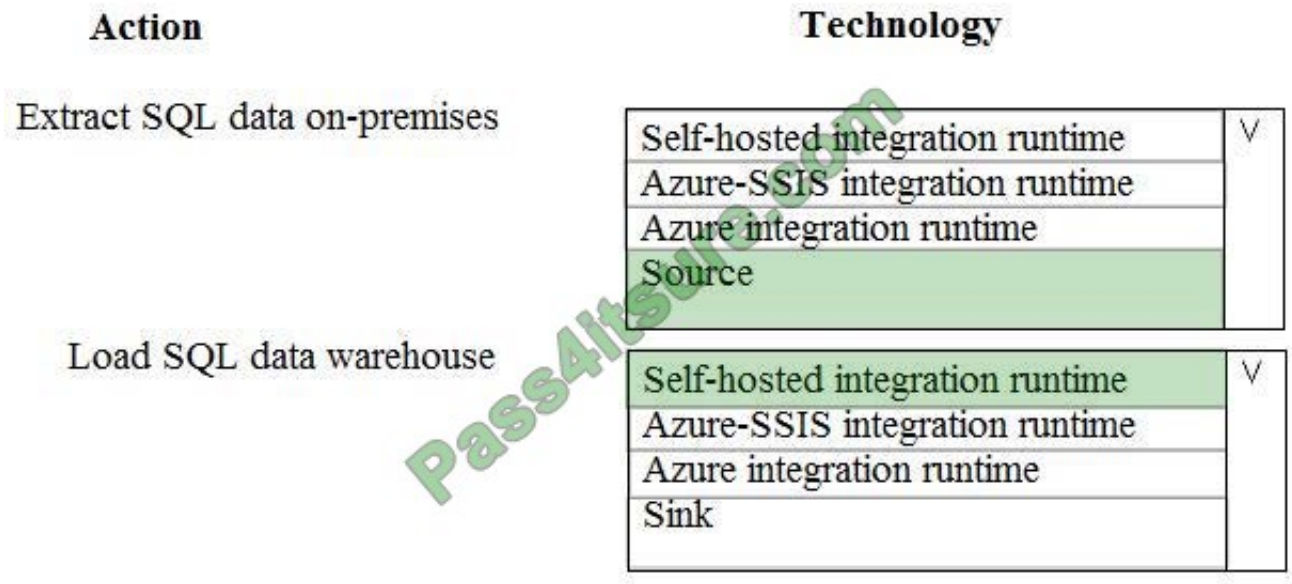

Box 1: Source

For Copy activity, it requires source and sink linked services to define the direction of data flow.

Copying between a cloud data source and a data source in private network: if either source or sink linked service points to a self-hosted IR, the copy activity is executed on that self-hosted Integration Runtime.

Box 2: Self-hosted integration runtime

A self-hosted integration runtime can run copy activities between a cloud data store and a data store in a private network, and it can dispatch transform activities against compute resources in an on-premises network or an Azure virtual

network. The installation of a self-hosted integration runtime needs on an on-premises machine or a virtual machine (VM) inside a private network.

References:

https://docs.microsoft.com/en-us/azure/data-factory/create-self-hosted-integration-runtime

QUESTION 4

You manage an enterprise data warehouse in Azure Synapse Analytics.

Users report slow performance when they run commonly used queries. Users do not report performance changes for infrequently used queries.

You need to monitor resource utilization to determine the source of the performance issues.

Which metric should you monitor?

A. Cache used percentage

B. Local tempdb percentage

C. DWU percentage

D. CPU percentage

E. Data IO percentage

Correct Answer: A

The Azure Synapse Analytics storage architecture automatically tiers your most frequently queried columnstore segments in a cache residing on NVMe based SSDs designed for Gen2 data warehouses. Greater performance is realized when your queries retrieve segments that are residing in the cache. You can monitor and troubleshoot slow query performance by determining whether your workload is optimally leveraging the Gen2 cache.

Note: As of November 2019, Azure SQL Data Warehouse is now Azure Synapse Analytics References: https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-how-to-monitor-cache

https://docs.microsoft.com/bs-latn-ba/azure/sql-data-warehouse/sql-data-warehouse-concept-resource-utilization-query-activity

QUESTION 5

You manage security for a database that supports a line of business application.

Private and personal data stored in the database must be protected and encrypted.

You need to configure the database to use Transparent Data Encryption (TDE).

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

Select and Place:

Actions	Answer Area
Create a database encryption key using a certificate generated with the master key.	
Create a certificate and then create the master key using a password.	
Set the context to the master database.	
Create a master key using a password.	
Set the context to the company database.	
Enable encryption.	

Correct Answer:

Actions	Answer Area
	Create a master key using a password.
	Create a certificate and then create the master key using a password.
Set the context to the master database.	Set the context to the company database.
	Create a database encryption key using a certificate generated with the master key.
	Enable encryption.

Step 1: Create a master key

Step 2: Create or obtain a certificate protected by the master key

Step 3: Set the context to the company database

Step 4: Create a database encryption key and protect it by the certificate

Step 5: Set the database to use encryption

Example code: USE master; GO CREATE MASTER KEY ENCRYPTION BY PASSWORD = \\'\\'; go CREATE CERTIFICATE MyServerCert WITH SUBJECT = \\'My DEK Certificate\\'; go USE AdventureWorks2012; GO CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = AES_128 ENCRYPTION BY SERVER CERTIFICATE MyServerCert; GO ALTER DATABASE AdventureWorks2012 SET ENCRYPTION ON; GO

References: https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption

QUESTION 6

You have a container named Sales in an Azure Cosmos DB database. Sales has 120 GB of data. Each entry in Sales has the following structure.

```
{
    OrderId: number,
    OrderDetailId: number,
    ProductName: string,
    other information that might vary…
}
```

The partition key is set to the OrderId attribute.

Users report that when they perform queries that retrieve data by ProductName, the queries take longer than expected to complete.

You need to reduce the amount of time it takes to execute the problematic queries.

Solution: You change the partition key to include ProductName.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

One option is to have a lookup collection "ProductName" for the mapping of "ProductName" to "OrderId".

References: https://azure.microsoft.com/sv-se/blog/azure-cosmos-db-partitioning-design-patterns-part-1/

QUESTION 7

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 questions 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 need to implement diagnostic logging for Data Warehouse monitoring.

Which log should you use?

A. RequestSteps

B. DmsWorkers

C. SqlRequests

D. ExecRequests

Correct Answer: C

Scenario:

The Azure SQL Data Warehouse cache must be monitored when the database is being used.

Metric	Description
A	Low cache hit %, high cache usage %
B	Low cache hit %, low cache usage %
C	High cache hit %, high cache usage %

References: https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-sql-requests-transact-sq

QUESTION 8

HOTSPOT

You have an Azure SQL database that contains a table named Customer. Customer contains the columns shown in the following table.

Customer_ID	Customer_Name	Customer_Phone
44531	John Smith	245-555-0173
44532	Tom Jones	245-505-3124
44533	Bill Tailor	245-689-4312

You plan to implement a dynamic data mask for the Customer_Phone column. The mask must meet the following requirements:

1.

The first six numerals of the customer phone numbers must be masked.

2.

The last four digits of the customer phone numbers must be visible.

3.

Hyphens must be preserved and displayed.

How should you configure the dynamic data mask? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

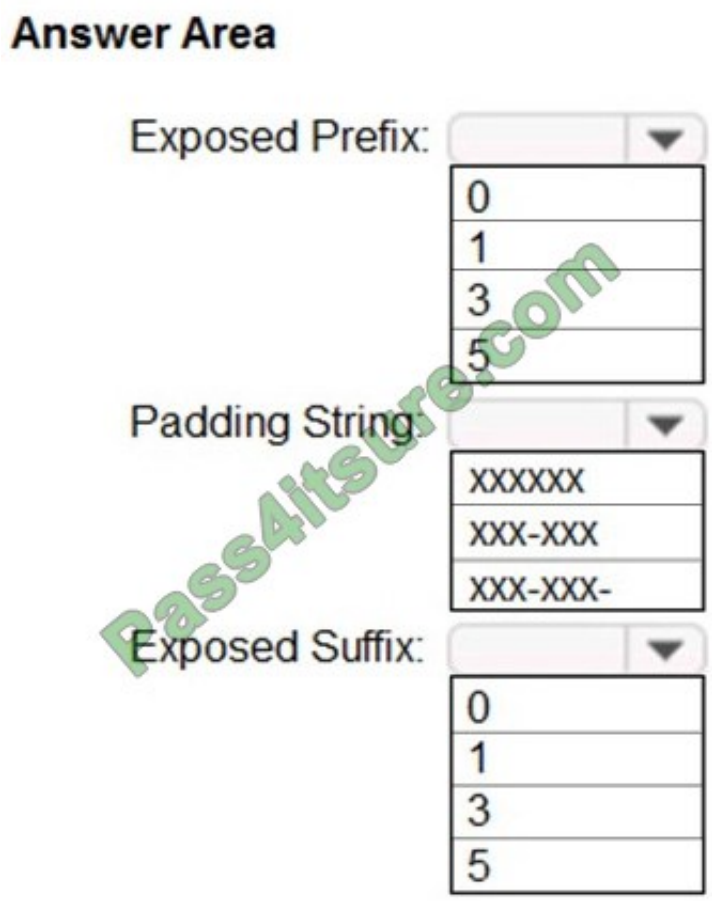

Correct Answer:

Answer Area

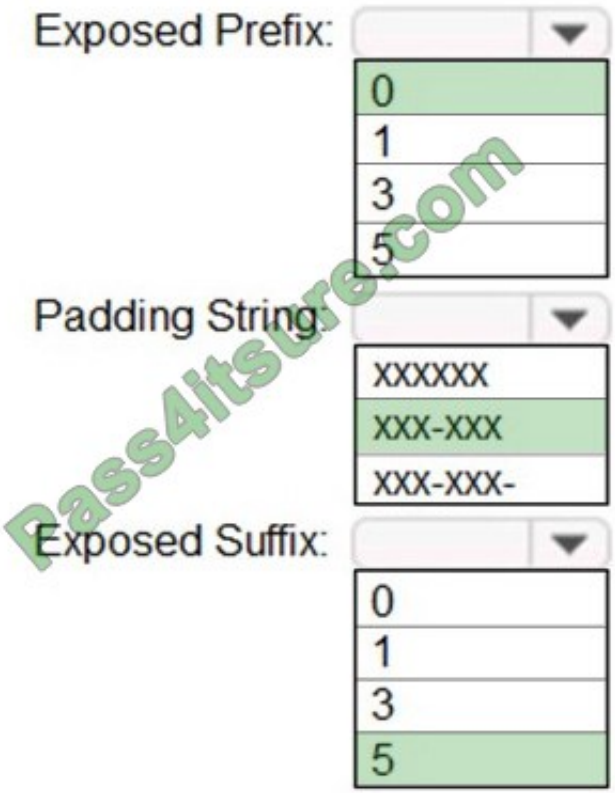

Exposed Prefix: 0

The first six digits must be masked. There is thus no exposed prefix.

Padding String: XXX-XXX

The first six digits must be masked and hyphens must be preserved.

Exposed Suffix: 5

The last 4 digits must be visible. There is not option for 4 but we can use 5 as it would include the hyphen before the last 4 digits.

Reference:

https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking?view=sql-server-ver15

QUESTION 9

A company uses Azure Data Lake Gen 1 Storage to store big data related to consumer behavior.

You need to implement logging.

Solution: Configure Azure Data Lake Storage diagnostics to store logs and metrics in a storage account.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: A

From the Azure Storage account that contains log data, open the Azure Storage account blade associated with Data Lake Storage Gen1 for logging, and then click Blobs. The Blob service blade lists two containers.

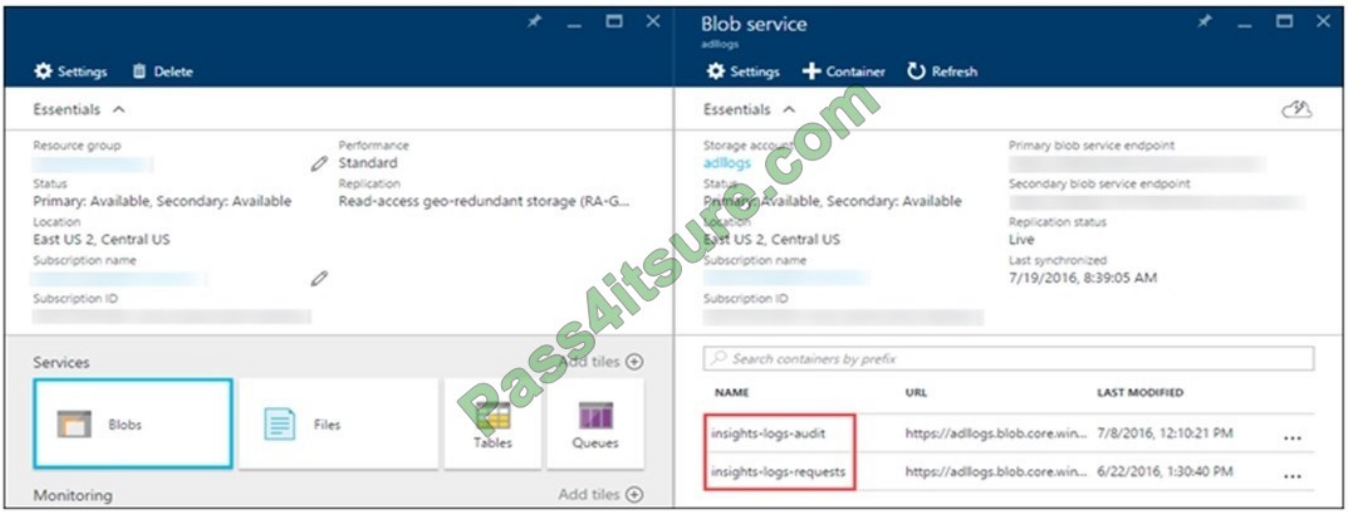

References: https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-diagnostic-logs

QUESTION 10

You are designing an enterprise data warehouse in Azure Synapse Analytics. You plan to load millions of rows of data into the data warehouse each day.

You must ensure that staging tables are optimized for data loading.

You need to design the staging tables.

What type of tables should you recommend?

A. Round-robin distributed table

B. Hash-distributed table

C. Replicated table

D. External table

Correct Answer: A

QUESTION 11

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 questions 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 need to configure data encryption for external applications.

Solution:

1.

Access the Always Encrypted Wizard in SQL Server Management Studio

2.

Select the column to be encrypted

3.

Set the encryption type to Deterministic

4.

Configure the master key to use the Windows Certificate Store

5.

Validate configuration results and deploy the solution

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

Use the Azure Key Vault, not the Windows Certificate Store, to store the master key.

Note: The Master Key Configuration page is where you set up your CMK (Column Master Key) and select the key store provider where the CMK will be stored. Currently, you can store a CMK in the Windows certificate store, Azure Key Vault, or a hardware security module (HSM).

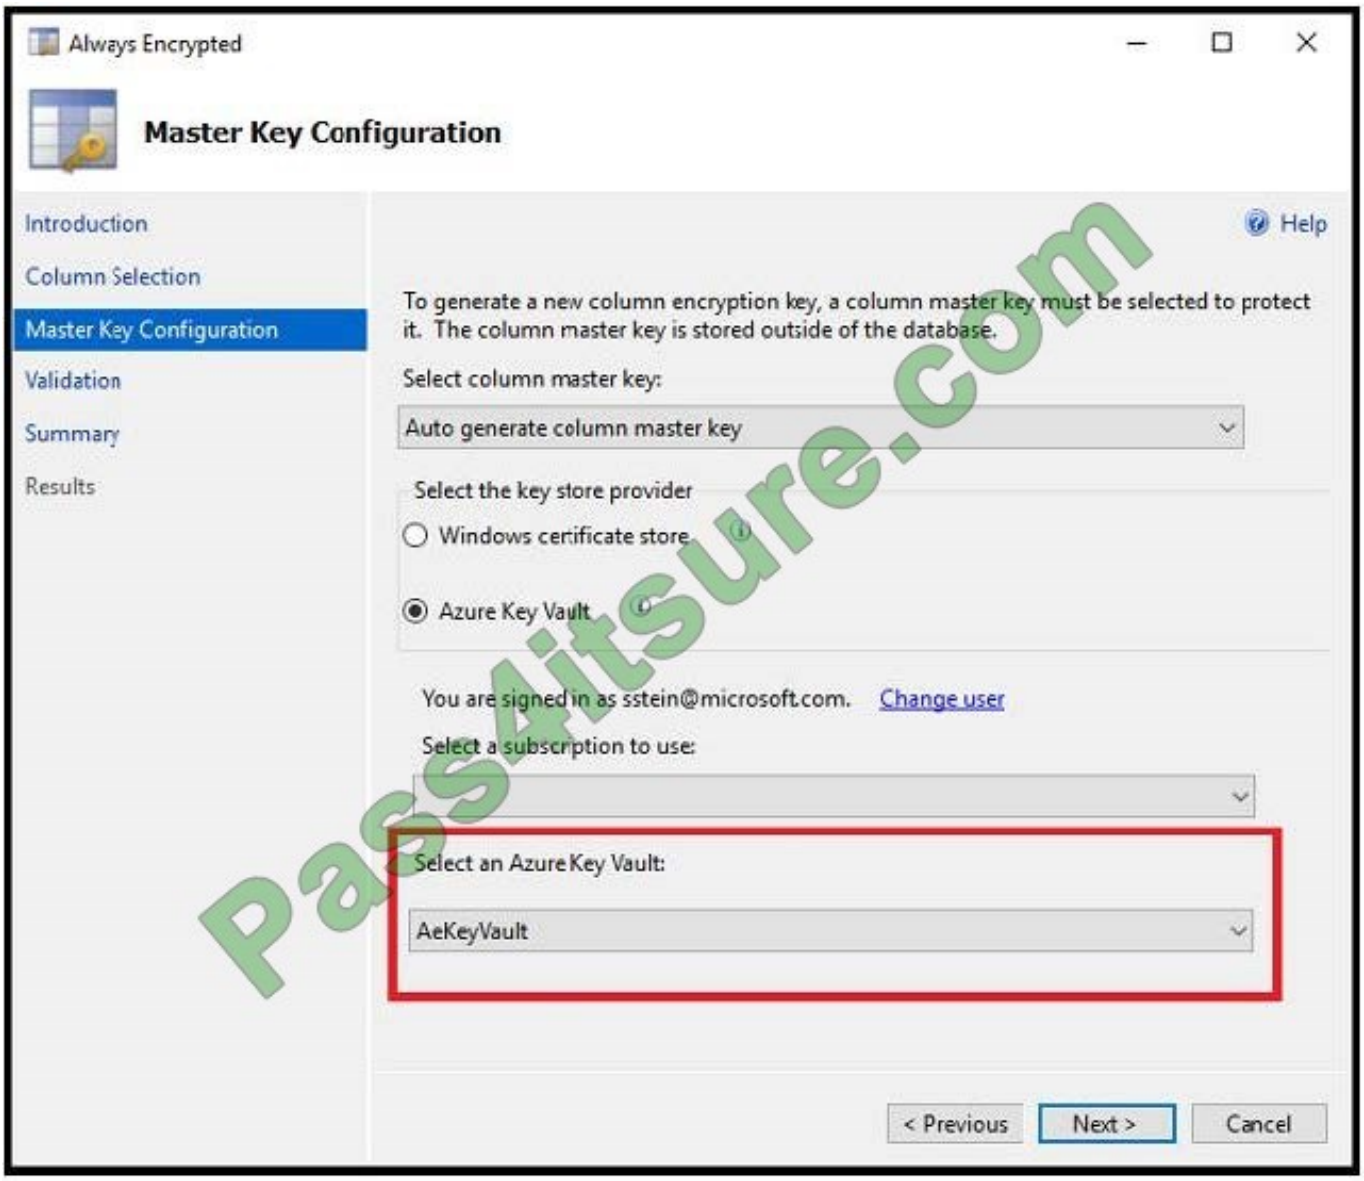

References: https://docs.microsoft.com/en-us/azure/sql-database/sql-database-always-encrypted-azure-key-vault

QUESTION 12

DRAG DROP

You have an Azure SQL database named DB1 in the Each US 2 region.

You need to build a secondary geo-replicated copy of DB1 in the West US region on a new server.

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

Select and Place:

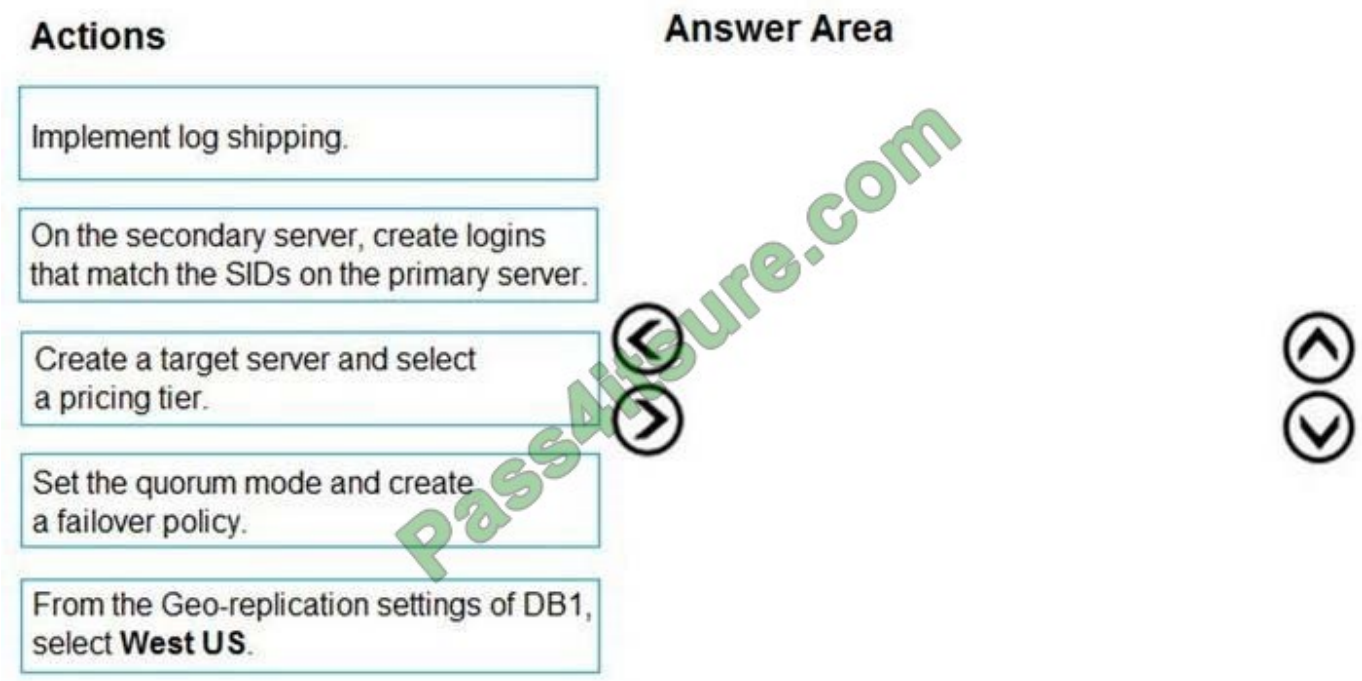

Correct Answer:

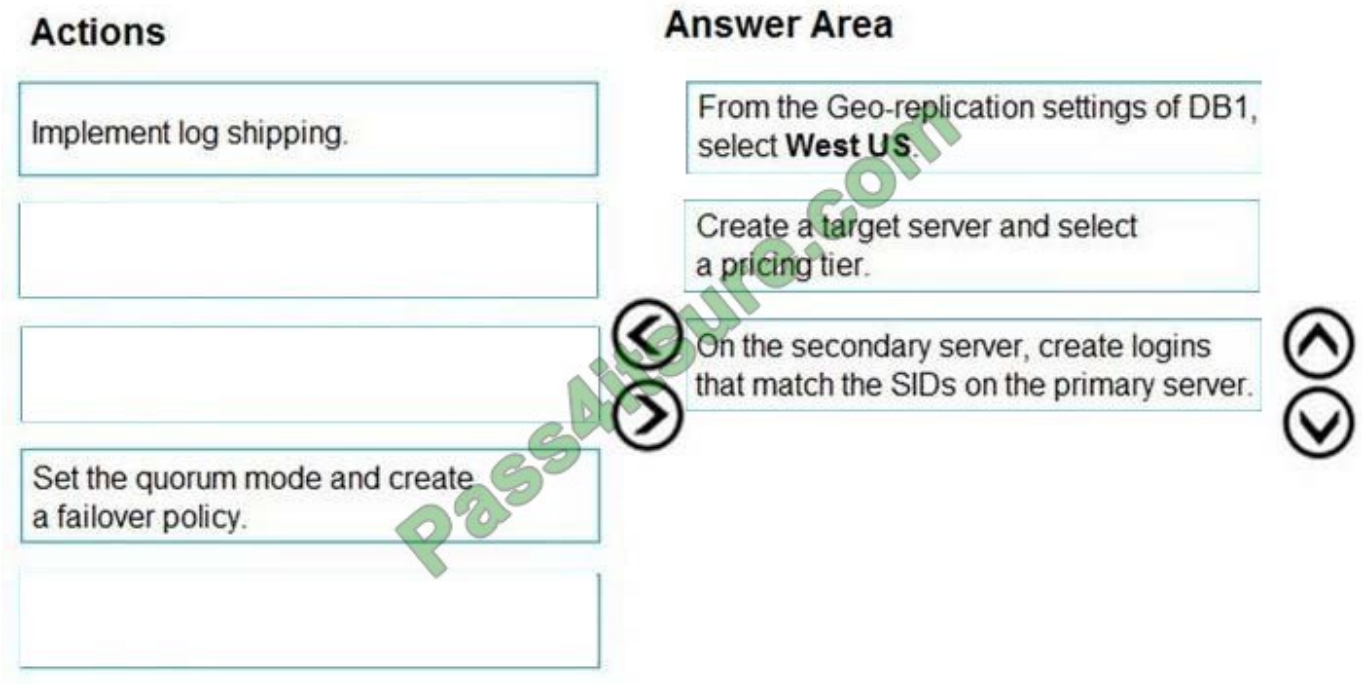

Step 1: From the Geo-replication settings of DB1, select West US

The following steps create a new secondary database in a geo-replication partnership.

1.

In the Azure portal, browse to the database that you want to set up for geo-replication.

2.

(Step 1) On the SQL database page, select geo-replication, and then select the region to create the secondary database.

3.

(Step 2-3) Select or configure the server and pricing tier for the secondary database.

Step 2: Create a target server and select a pricing tier

Step 3: On the secondary server, create logins that match the SIDs on the primary server.

Incorrect Answers:

Not log shipping: Replication is used.

References:

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-active-geo-replication-portal