Microsoft

(DP-300)

Administering Relational Databases on Microsoft Azure

Total: **361 Questions**

Link:

Question: 1

You have 20 Azure SQL databases provisioned by using the vCore purchasing model.
You plan to create an Azure SQL Database elastic pool and add the 20 databases.
Which three metrics should you use to size the elastic pool to meet the demands of your workload? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. total size of all the databases
B. geo-replication support
C. number of concurrently peaking databases * peak CPU utilization per database
D. maximum number of concurrent sessions for all the databases
E. total number of databases * average CPU utilization per database

Answer: ACE

Explanation:
CE: Estimate the vCores needed for the pool as follows:
For vCore-based purchasing model: MAX(<Total number of DBs X average vCore utilization per DB>, <Number of concurrently peaking DBs X Peak vCore utilization per DB)
A: Estimate the storage space needed for the pool by adding the number of bytes needed for all the databases in the pool.

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

Question: 2

DRAG DROP -
You have SQL Server 2019 on an Azure virtual machine that contains an SSISDB database.
A recent failure causes the master database to be lost.
You discover that all Microsoft SQL Server integration Services (SSIS) packages fail to run on the virtual machine. Which four actions should you perform in sequence to resolve the issue? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct.
Select and Place:

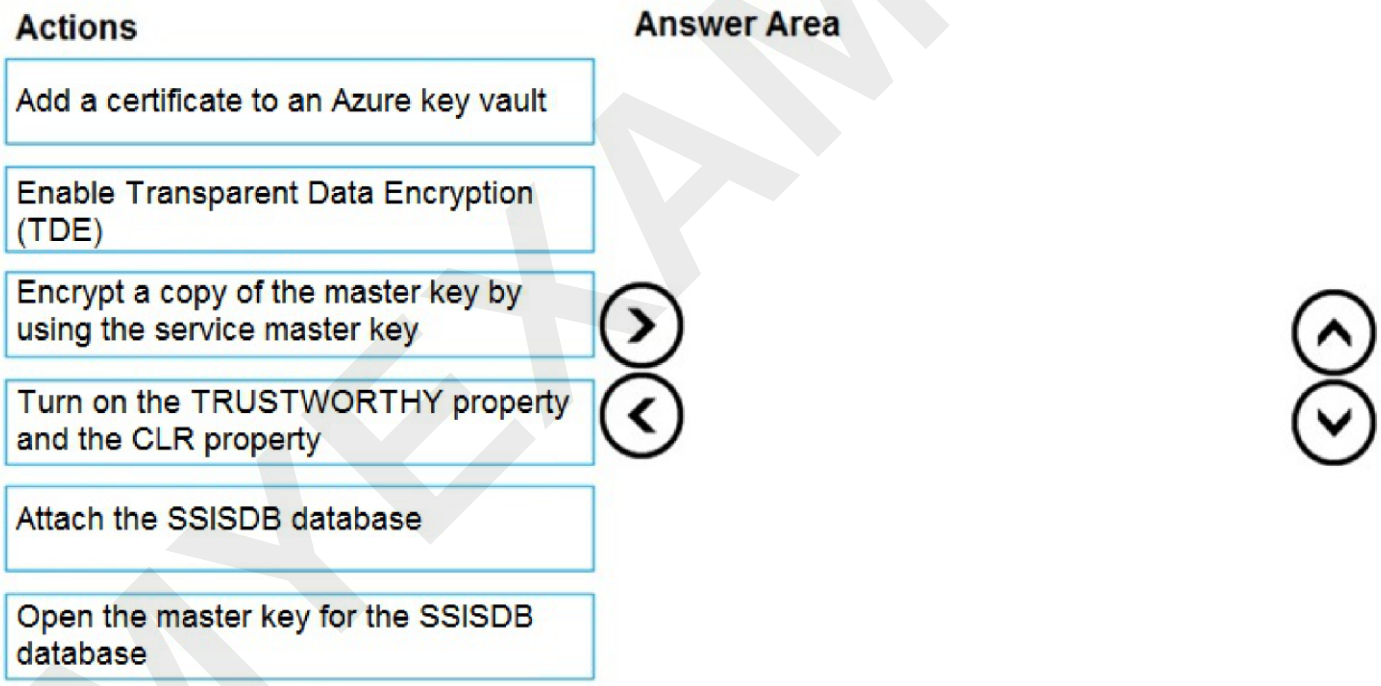

Answer:

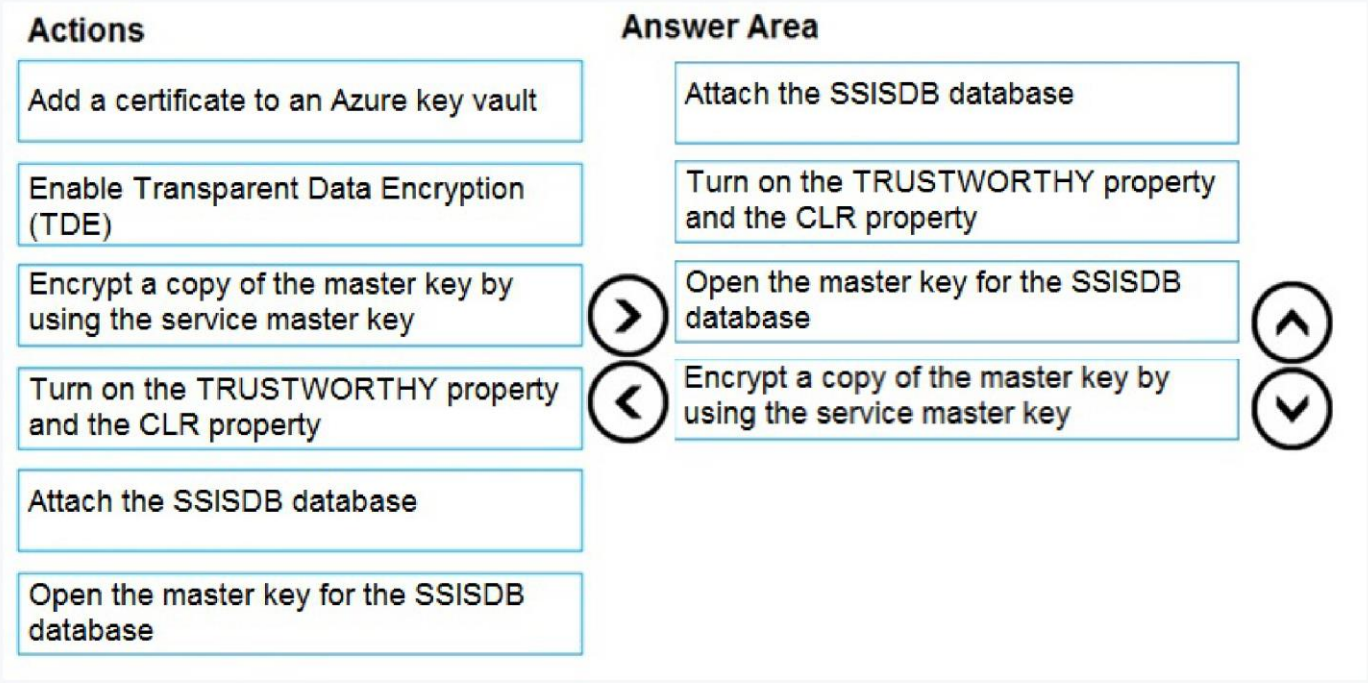

Explanation:
Step 1: Attach the SSISDB database
Step 2: Turn on the TRUSTWORTHY property and the CLR property
If you are restoring the SSISDB database to an SQL Server instance where the SSISDB catalog was never created, enable common language runtime (clr)
Step 3: Open the master key for the SSISDB database
Restore the master key by this method if you have the original password that was used to create SSISDB.

open master key decryption by password = 'LS1Setup!' --'Password used when creating SSISDB' Alter Master Key Add encryption by Service Master Key
Step 4: Encrypt a copy of the master key by using the service master key

Reference:
https://docs.microsoft.com/en-us/sql/integration-services/catalog/ssis-catalog

Question: 3

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

Name	Data type
SalesID	Int
Product	Int
Total Number	Numeric(8,4)
Tax Number	Numeric(8,4)
SalesRep	Varchar(50)

FactSales has 6 billion rows and is loaded nightly by using a batch process. You must provide the greatest reduction in space for the database and maximize performance.
Which type of compression provides the greatest space reduction for the database?

A. page compression
B. row compression
C. columnstore compression
D. columnstore archival compression

Answer: D

Explanation:

Columnstore tables and indexes are always stored with columnstore compression. You can further reduce the size of columnstore data by configuring an additional compression called archival compression.

Note: Columnstore " The columnstore index is also logically organized as a table with rows and columns, but the data is physically stored in a column-wise data format.

Incorrect Answers:

B: Rowstore " The rowstore index is the traditional style that has been around since the initial release of SQL Server.

For rowstore tables and indexes, use the data compression feature to help reduce the size of the database. C: Columnstore compression is less compressed compared to columnstore archival compression.

Reference:

https://docs.microsoft.com/en-us/sql/relational-databases/data-compression/data-compression

Question: 4 You have a Microsoft SQL Server 2019 database named DB1 that uses the following database-level and instance-level features.

- Clustered columnstore indexes
- Automatic tuning
- Change tracking
- PolyBase

You plan to migrate DB1 to an Azure SQL database.
What feature should be removed or replaced before DB1 can be migrated?

A. Clustered columnstore indexes
B. PolyBase
C. Change tracking
D. Automatic tuning

Answer: B

Explanation:

This table lists the key features for PolyBase and the products in which they're available.

Feature	SQL Server (**Beginning with 2016**)	Azure SQL Database	Azure Synapse Analytics	Parallel Data Warehouse
Query Hadoop data with Transact-SQL	Yes	No	No	Yes
Import data from Hadoop	Yes	No	No	Yes
Export data to Hadoop	Yes	No	No	Yes
Query, import from, export to Azure HDInsight	No	No	No	No
Push down query computations to Hadoop	Yes	No	No	Yes
Import data from Azure Blob storage	Yes	Yes*	Yes	Yes
Export data to Azure Blob storage	Yes	No	Yes	Yes
Import data from Azure Data Lake Store	No	No	Yes	No
Export data to Azure Data Lake Store	No	No	Yes	No
Run PolyBase queries from Microsoft BI tools	Yes	No	Yes	Yes

Incorrect Answers:
C: Change tracking is a lightweight solution that provides an efficient change tracking mechanism for applications. It applies to both Azure SQL Database and SQL
Server.

D: Azure SQL Database and Azure SQL Managed Instance automatic tuning provides peak performance and stable workloads through continuous performance tuning based on AI and machine learning.

Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-versioned-feature-summary

Question: 5

You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4-TB database named DB1.
You plan to migrate DB1 to an Azure SQL Database managed instance.
What should you use to minimize downtime and data loss during the migration?

A. distributed availability groups
B. database mirroring
C. Always On Availability Group
D. Azure Database Migration Service

Answer: D

Explanation:

Azure Database Migration Service can do online migrations with minimal downtime. When you migrate databases to Azure by using Azure Database Migration Service, you can do an offline or an online migration. With an offline migration, application downtime starts when the migration starts. With an online migration, downtime is limited to the time to cut over at the end of migration. We suggest that you test an offline

migration to determine whether the downtime is acceptable; if not, do an online migration.

Reference:

https://docs.microsoft.com/en-us/azure/dms/dms-overview

Question: 6

HOTSPOT -
You have an on-premises Microsoft SQL Server 2016 server named Server1 that contains a database named DB1.
You need to perform an online migration of DB1 to an Azure SQL Database managed instance by using Azure Database Migration Service.
How should you configure the backup of DB1? To answer, select the appropriate options in the answer area.
NOTE:
Each correct selection is worth one point.
Hot Area:

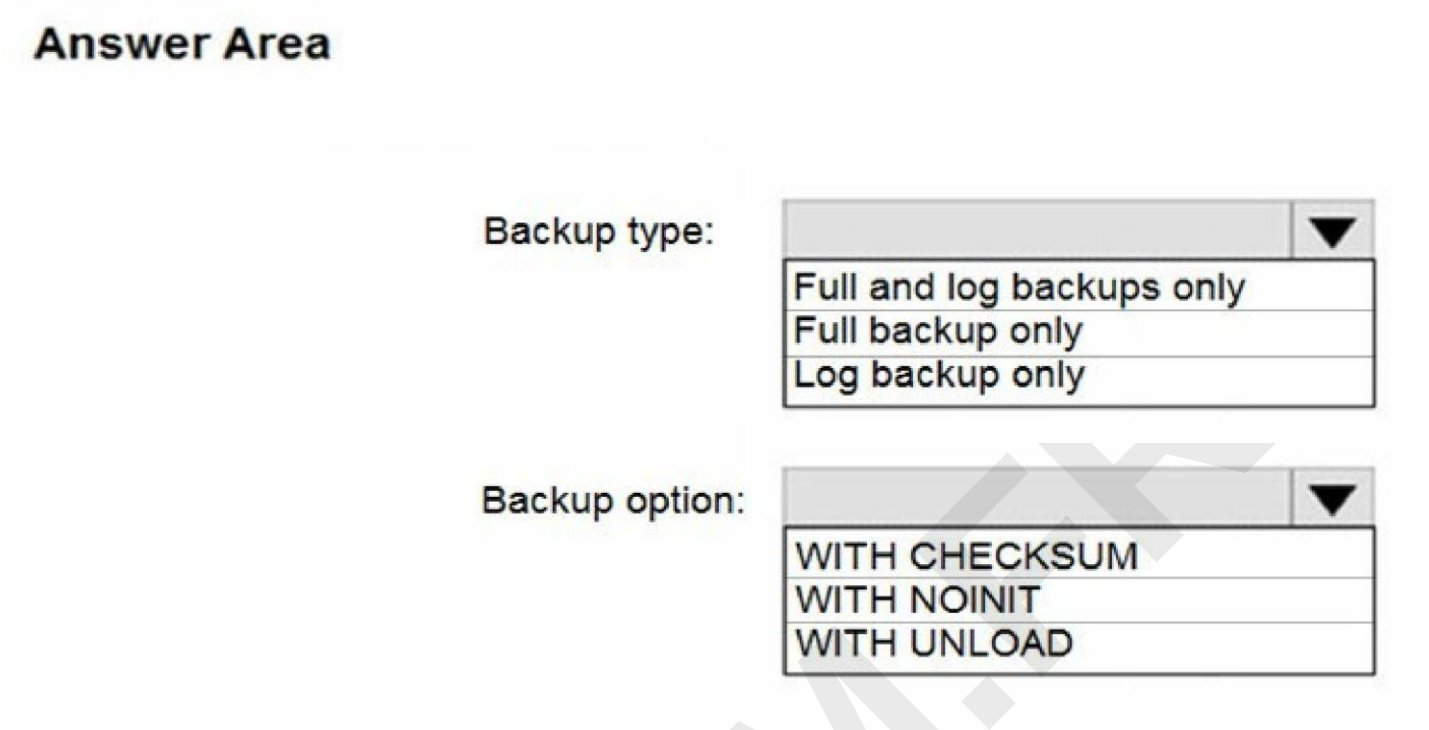

Answer:

Answer Area

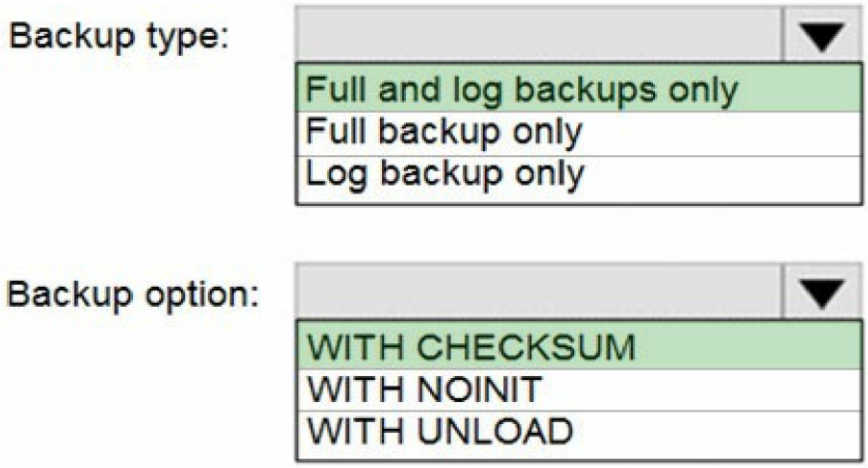

Explanation:
Box 1: Full and log backups only
Make sure to take every backup on a separate backup media (backup files). Azure Database Migration Service doesn't support backups that are appended to a single backup file. Take full backup and log backups to separate backup files.

Box 2: WITH CHECKSUM -
Azure Database Migration Service uses the backup and restore method to migrate your on-premises databases to SQL Managed Instance. Azure Database
Migration Service only supports backups created using checksum.

Incorrect Answers:
NOINIT Indicates that the backup set is appended to the specified media set, preserving existing backup sets.
If a media password is defined for the media set, the password must be supplied. NOINIT is the default.

UNLOAD -
Specifies that the tape is automatically rewound and unloaded when the backup is finished. UNLOAD is the default when a session begins.

Reference:
https://docs.microsoft.com/en-us/azure/dms/known-issues-azure-sql-db-managed-instance-online

Question: 7

DRAG DROP -
You have a resource group named App1Dev that contains an Azure SQL Database server named DevServer1.
DevServer1 contains an Azure SQL database named DB1. The schema and permissions for DB1 are saved in a Microsoft SQL Server Data Tools (SSDT) database project.
You need to populate a new resource group named App1Test with the DB1 database and an Azure SQL Server named TestServer1. The resources in App1Test must have the same configurations as the resources in App1Dev.
Which four 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:

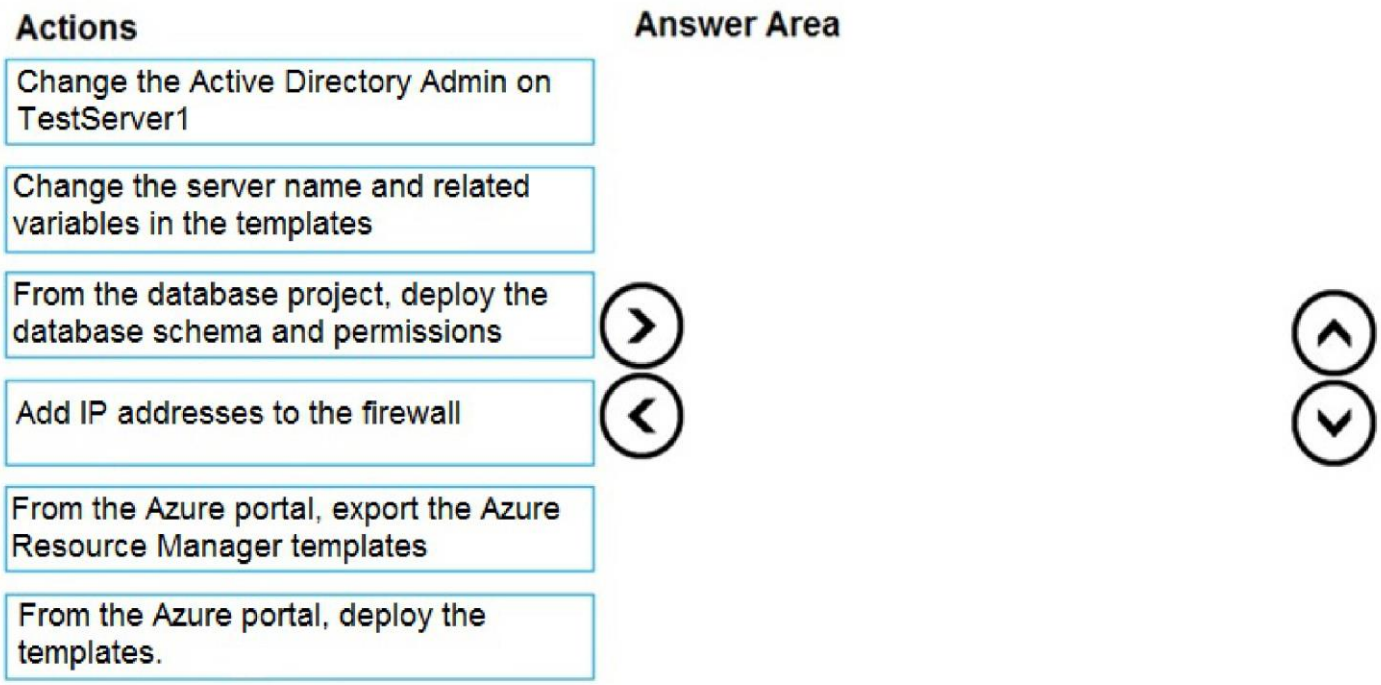

Answer:

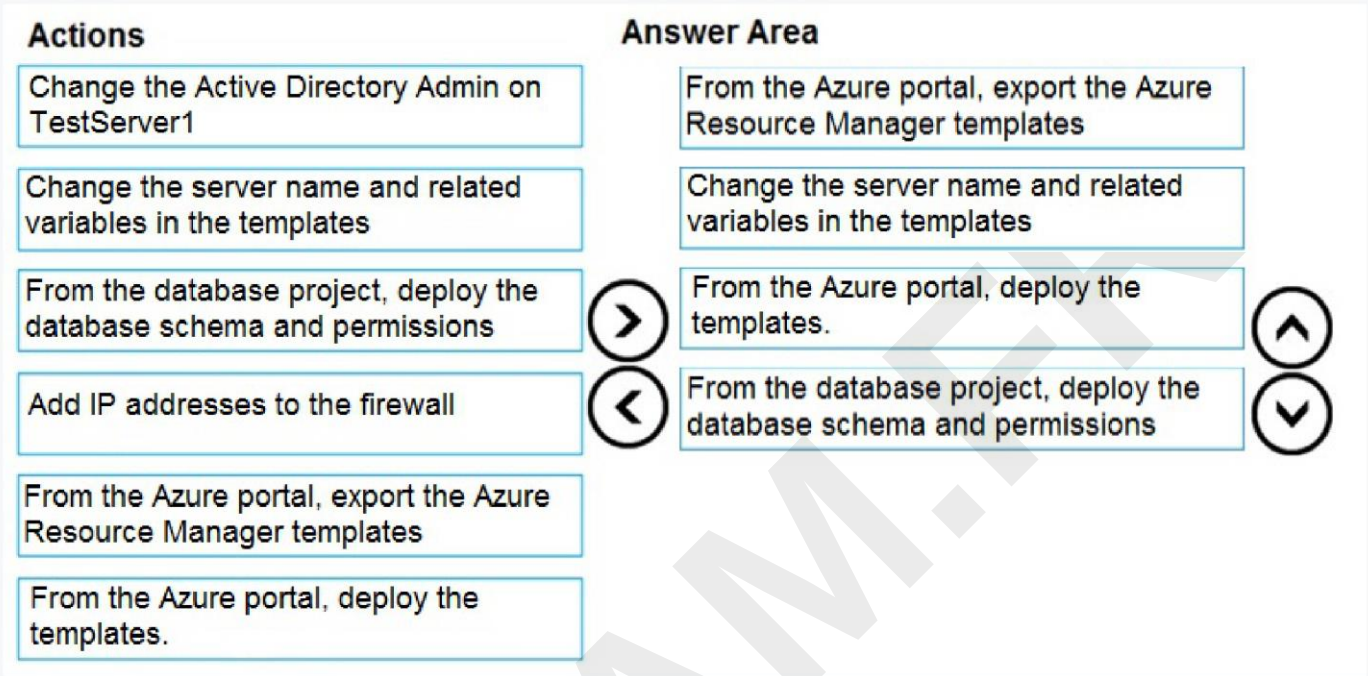

Explanation:

1. From the Azure Portal, export the Azure Resource Manager templates. 2.

Change the server name and related variables in the templates.

3. From the Azure portal, deploy the templates.

4. From the database project, deply the database schema and permissions.

Question: 8 HOTSPOT

-

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and an Azure Data Lake Storage Gen2 account named Account1.

You plan to access the files in Account1 by using an external table.
You need to create a data source in Pool1 that you can reference when you create the external table.
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

```
CREATE EXTERNAL DATA SOURCE source1

WITH

      ( LOCATION = 'https://account1.  [ ▼ ]  .core.windows.net',
                                        blob
                                        dfs
                                        table

        [ ▼ ]
        PUSHDOWN = ON
        TYPE = BLOB_STORAGE
        TYPE = HADOOP

        )
```

Answer:

Answer Area

```
CREATE EXTERNAL DATA SOURCE source1

WITH

      ( LOCATION = 'https://account1.  [ ▼ ]  .core.windows.net',
                                        blob
                                        dfs
                                        table

        [ ▼ ]
        PUSHDOWN = ON
        TYPE = BLOB_STORAGE
        TYPE = HADOOP

        )
```


Explanation:
Box 1: dfs -
For Azure Data Lake Store Gen 2 used the following syntax:
http[s] <storage_account>.dfs.core.windows.net/<container>/subfolders
Incorrect:
Not blob: blob is used for Azure Blob Storage. Syntax:
http[s] <storage_account>.blob.core.windows.net/<container>/subfolders

Box 2: TYPE = HADOOP -
Syntax for CREATE EXTERNAL DATA SOURCE.
External data sources with TYPE=HADOOP are available only in dedicated SQL pools.
CREATE EXTERNAL DATA SOURCE <data_source_name>

WITH -
(LOCATION = '<prefix>://<path>'
[, CREDENTIAL = <database scoped credential>]
, TYPE = HADOOP
)
[;]

Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables

Question: 9

HOTSPOT -
You plan to develop a dataset named Purchases by using Azure Databricks. Purchases will contain the following columns:
- ProductID
- ItemPrice
- LineTotal
- Quantity
- StoreID
- Minute
- Month
- Hour
- Year
- Day

You need to store the data to support hourly incremental load pipelines that will vary for each StoreID. The solution must minimize storage costs.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

```
df.write
```

.bucketBy	(“*”)
.partitionBy	(“StoreID”, “Hour”)
.range	(“StoreID”, “Year”, “Month”, “Day”, “Hour”)
.sortBy	(“Year”, “Month”, “Day”, “Hour” “StoreID”)

```
.mode(“append”)
```

.csv(“/Purchases”)
.json(“/Purchases”)
.parquet(“/Purchases”)
.saveAsTable(“/Purchases”)

Answer:

Answer Area

```
df.write
```

.bucketBy	(“*”)
.partitionBy	(“StoreID”, “Hour”)
.range	(“StoreID”, “Year”, “Month”, “Day”, “Hour”)
.sortBy	(“Year”, “Month”, “Day”, “Hour” “StoreID”)

```
.mode(“append”)
```

.csv(“/Purchases”)
.json(“/Purchases”)
.parquet(“/Purchases”)
.saveAsTable(“/Purchases”)

Explanation:
Box 1: .partitionBy -
Example:
df.write.partitionBy("y","m","d")
.mode(SaveMode.Append)
.parquet("/data/hive/warehouse/db_name.db/" + tableName)

Box 2: ("Year","Month","Day","Hour","StoreID")
Box 3: .parquet("/Purchases")

Reference:
https://intellipaat.com/community/11744/how-to-partition-and-write-dataframe-in-spark-without-deleting-partitions-with-no-new-data

Question: 10

You are designing a streaming data solution that will ingest variable volumes of data. You need to ensure that you can change the partition count after creation.
Which service should you use to ingest the data?

A. Azure Event Hubs Standard
B. Azure Stream Analytics
C. Azure Data Factory
D. Azure Event Hubs Dedicated

Answer: D

Explanation:
The partition count for an event hub in a dedicated Event Hubs cluster can be increased after the event hub has been created.
Incorrect Answers:
A: For Azure Event standard hubs, the partition count isn't changeable, so you should consider long-term scale when setting partition count.

Reference:
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features#partitions

Question: 11

HOTSPOT -
You are building a database in an Azure Synapse Analytics serverless SQL pool. You have data stored in Parquet files in an Azure Data Lake Storage Gen2 container. Records are structured as shown in the following sample.

```
{
    "id":123,
    "address_housenumber": "19c",
    "address_line1": "Memory Lane",
    "applicant1_name": "Jane",
    "applicant2_name": "Dev"
}
```

The records contain two applicants at most.
You need to build a table that includes only the address fields.
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

[▼] applications

- CREATE EXTERNAL TABLE
- CREATE TABLE
- CREATE VIEW

```
WITH (
     LOCATION = 'applications/',
     DATA_SOURCE = applications_ds,
     FILE_FORMAT = applications_file_format
)
AS
SELECT id, [address_housenumber] as addressnumber, [address_line1]
as addressline1
FROM
```

[▼] (BULK 'https://contoso1.dfs.core.windows.net/applications/year=*/*.parquet',

- CROSS APPLY
- OPENJSON
- OPENROWSET

```
     FORMAT = 'PARQUET') AS [r]
GO
```

Answer:

Answer Area

[▼] applications

- CREATE EXTERNAL TABLE
- CREATE TABLE
- CREATE VIEW

```
WITH (
     LOCATION = 'applications/',
     DATA_SOURCE = applications_ds,
     FILE_FORMAT = applications_file_format
)
AS
SELECT id, [address_housenumber] as addressnumber, [address_line1]
as addressline1
FROM
```

[▼] (BULK 'https://contoso1.dfs.core.windows.net/applications/year=*/*.parquet',

- CROSS APPLY
- OPENJSON
- OPENROWSET

```
     FORMAT = 'PARQUET') AS [r]
GO
```


Explanation:

Box 1: CREATE EXTERNAL TABLE -

An external table points to data located in Hadoop, Azure Storage blob, or Azure Data Lake Storage. External tables are used to read data from files or write data to files in Azure Storage. With Synapse SQL, you can use external tables to read external data using dedicated SQL pool or serverless SQL pool.

Syntax:
CREATE EXTERNAL TABLE database_name.schema_name.table_name | schema_name.table_name | table_name
(<column_definition> [,...n])
WITH (
LOCATION = 'folder_or_filepath',
DATA_SOURCE = external_data_source_name,
FILE_FORMAT = external_file_format_name

Box 2. OPENROWSET -
When using serverless SQL pool, CETAS is used to create an external table and export query results to Azure Storage Blob or Azure Data Lake Storage Gen2.

Example:

AS -
SELECT decennialTime, stateName, SUM(population) AS population

FROM -
OPENROWSET(BULK 'https://azureopendatastorage.blob.core.windows.net/censusdatacontainer/release/us_population_county/year=*/*
FORMAT='PARQUET') AS [r]
GROUP BY decennialTime, stateName

GO -

Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables

Question: 12

You have an Azure Synapse Analytics Apache Spark pool named Pool1.
You plan to load JSON files from an Azure Data Lake Storage Gen2 container into the tables in Pool1. The structure and data types vary by file.
You need to load the files into the tables. The solution must maintain the source data types.
What should you do?

A. Load the data by using PySpark.

B. Load the data by using the OPENROWSET Transact-SQL command in an Azure Synapse Analytics serverless SQL pool.

C. Use a Get Metadata activity in Azure Data Factory.

D. Use a Conditional Split transformation in an Azure Synapse data flow.

Answer: A

Explanation:
Synapse notebooks support four Apache Spark languages:
PySpark (Python)
Spark (Scala)

Spark SQL -

.NET Spark (C#)
Note: Bring data to a notebook.

You can load data from Azure Blob Storage, Azure Data Lake Store Gen 2, and SQL pool as shown in the code samples below.

Read a CSV from Azure Data Lake Store Gen2 as a Spark DataFrame. from pyspark.sql import SparkSession from pyspark.sql.types import * account_name = "Your account name" container_name = "Your container name" relative_path = "Your path" adls_path = 'abfss://%%s.dfs.core.windows.net/%s' % (container_name, account_name, relative_path) df1 = spark.read.option('header', 'true') \
.option('delimiter', ',') \
.csv(adls_path + '/Testfile.csv')

Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/spark/apache-spark-development-using-notebook s

Question: 13

You are designing a date dimension table in an Azure Synapse Analytics dedicated SQL pool. The date dimension table will be used by all the fact tables.
Which distribution type should you recommend to minimize data movement?

A. HASH
B. REPLICATE
C. ROUND_ROBIN

Answer: B

Explanation:
A replicated table has a full copy of the table available on every Compute node. Queries run fast on replicated tables since joins on replicated tables don't require data movement. Replication requires extra storage, though, and isn't practical for large tables.

Incorrect Answers:
C: A round-robin distributed table distributes table rows evenly across all distributions. The assignment of rows to distributions is random. Unlike hash-distributed tables, rows with equal values are not guaranteed to be assigned to the same distribution.

As a result, the system sometimes needs to invoke a data movement operation to better organize your data before it can resolve a query.

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

Question: 14

HOTSPOT -
From a website analytics system, you receive data extracts about user interactions such as downloads, link clicks, form submissions, and video plays.
The data contains the following columns:

Name	Sample value
Date	15 Jan 2021
EventCategory	Videos
EventAction	Play
EventLabel	Contoso Promotional
ChannelGrouping	Social
TotalEvents	150
UniqueEvents	120
SessionsWithEvents	99

You need to design a star schema to support analytical queries of the data. The star schema will contain four tables including a date dimension.
To which table should you add each column? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

EventCategory: DimChannel / DimDate / DimEvent / FactEvents

ChannelGrouping: DimChannel / DimDate / DimEvent / FactEvents

TotalEvents: DimChannel / DimDate / DimEvent / FactEvents

Answer:

Answer Area

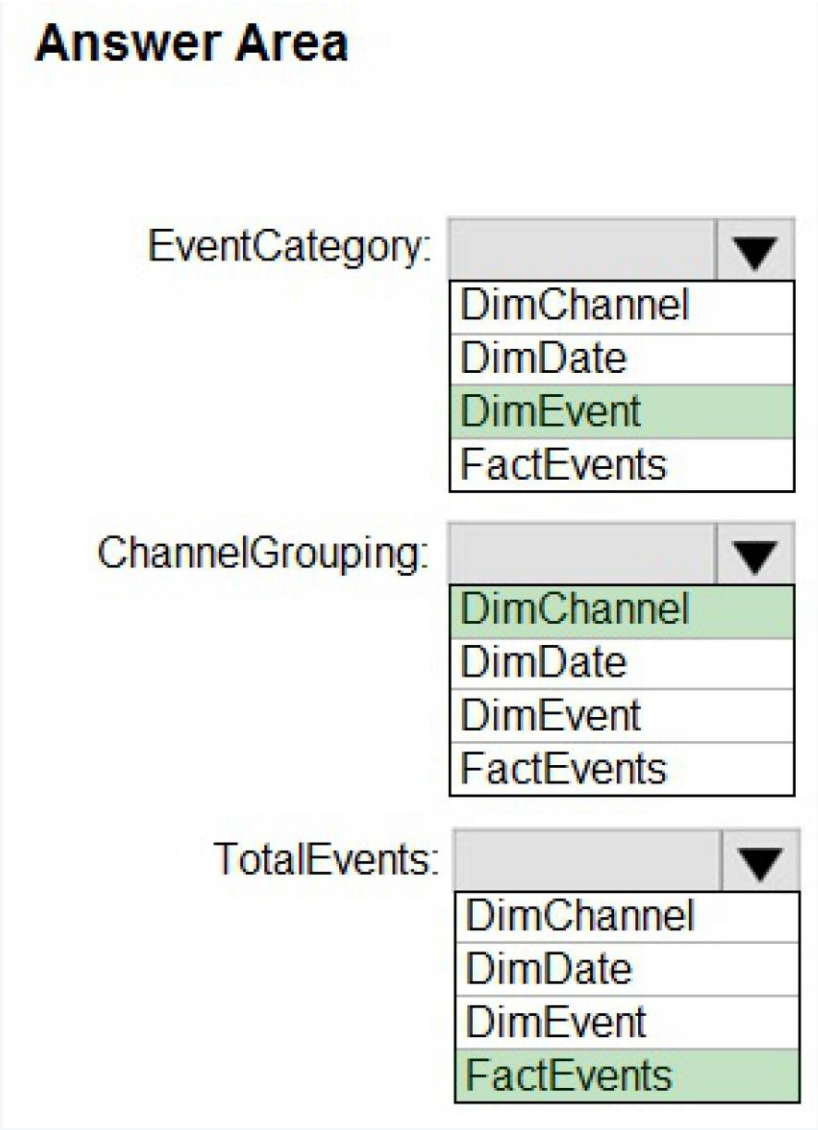

Explanation:
Box 1: DimEvent -

Box 2: DimChannel -
Dimension tables describe business entities " the things you model. Entities can include products, people, places, and concepts including time itself. The most consistent table you'll find in a star schema is a date dimension table. A dimension table contains a key column (or columns) that acts as a unique identifier, and descriptive columns.

Box 3: FactEvents -
Fact tables store observations or events, and can be sales orders, stock balances, exchange rates, temperatures, etc.

Reference:
https://docs.microsoft.com/en-us/power-bi/guidance/star-schema

Question: 15

DRAG DROP -
You plan to create a table in an Azure Synapse Analytics dedicated SQL pool.
Data in the table will be retained for five years. Once a year, data that is older than five years will be deleted. You need to ensure that the data is distributed evenly across partitions. The solutions must minimize the amount of time required to delete old data.
How should you complete the Transact-SQL statement? To answer, drag the appropriate values to the correct targets. Each value 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.
Select and Place:

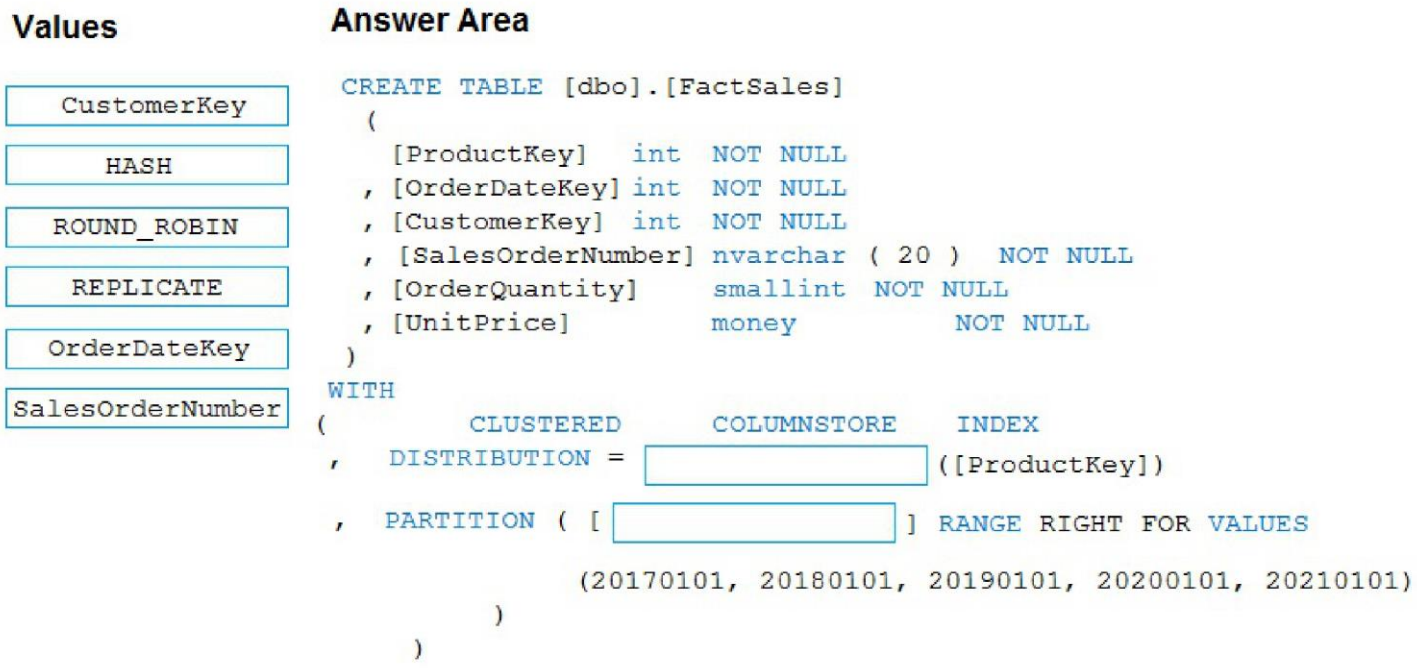

Answer:

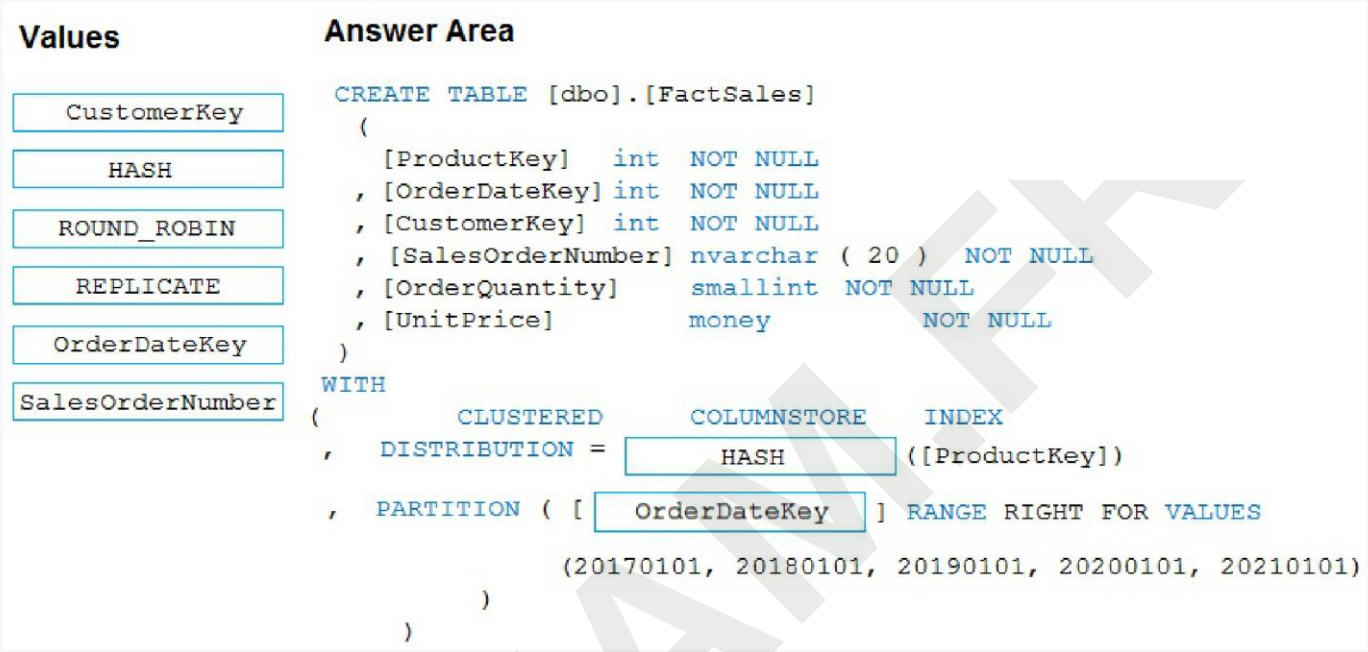

Explanation:
Box 1: HASH -

Box 2: OrderDateKey -
In most cases, table partitions are created on a date column.

A way to eliminate rollbacks is to use Metadata Only operations like partition switching for data management. For example, rather than execute a DELETE statement to delete all rows in a table where the order_date was in October of 2001, you could partition your data early. Then you can switch out the partition with data for an empty partition from another table.

Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/create-table-azure-sql-data-warehouse https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/best-practices-dedicated-sql-pool

Question: 16

You have an Azure Synapse Analytics workspace named WS1 that contains an Apache Spark pool named Pool1. You plan to create a database named DB1 in Pool1.
You need to ensure that when tables are created in DB1, the tables are available automatically as external tables to the built-in serverless SQL pool.
Which format should you use for the tables in DB1?

A. JSON
B. CSV
C. Parquet
D. ORC

Answer: C

Explanation:
Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools.

For each Spark external table based on Parquet and located in Azure Storage, an external table is created in a serverless SQL pool database. As such, you can shut down your Spark pools and still query Spark external tables from serverless SQL pool.

Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-storage-files-spark-tables

Question: 17

You are designing an anomaly detection solution for streaming data from an Azure IoT hub. The solution must meet the following requirements:
- Send the output to an Azure Synapse.
- Identify spikes and dips in time series data.
- Minimize development and configuration effort.

Which should you include in the solution?

A. Azure SQL Database
B. Azure Databricks
C. Azure Stream Analytics

Answer: C

Explanation:
Anomalies can be identified by routing data via IoT Hub to a built-in ML model in Azure Stream Analytics

Reference:
https://docs.microsoft.com/en-us/learn/modules/data-anomaly-detection-using-azure-iot-hub/ https://docs.microsoft.com/en-us/azure/stream-analytics/azure-synapse-analytics-output

Question: 18

You are creating a new notebook in Azure Databricks that will support R as the primary language but will also support Scala and SQL.
Which switch should you use to switch between languages?

A. \\[<language>]

B. %<language>
C. \\[<language>]
D. @<language>

Answer: B

Explanation:
You can override the default language by specifying the language magic command %<language> at the beginning of a cell. The supported magic commands are:
%python, %r, %scala, and %sql.

Reference:
https://docs.microsoft.com/en-us/azure/databricks/notebooks/notebooks-use

Question: 19

DRAG DROP -
You are creating a managed data warehouse solution on Microsoft Azure.
You must use PolyBase to retrieve data from Azure Blob storage that resides in parquet format and load the data into a large table called FactSalesOrderDetails.
You need to configure Azure Synapse Analytics to receive the data.
Which four 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:

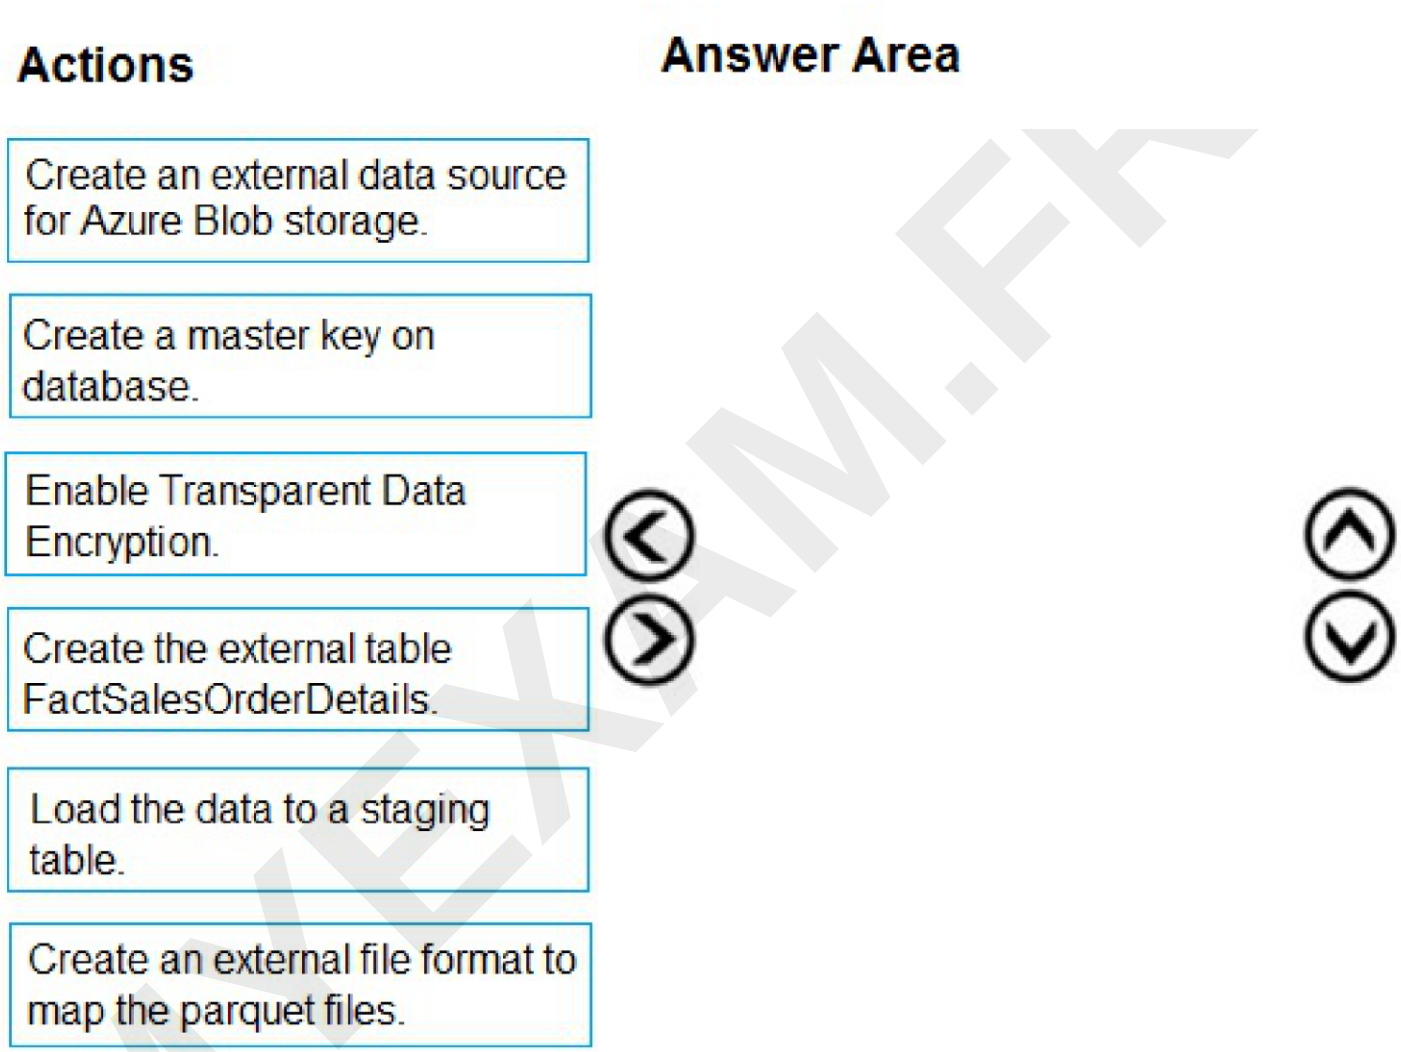

Answer:

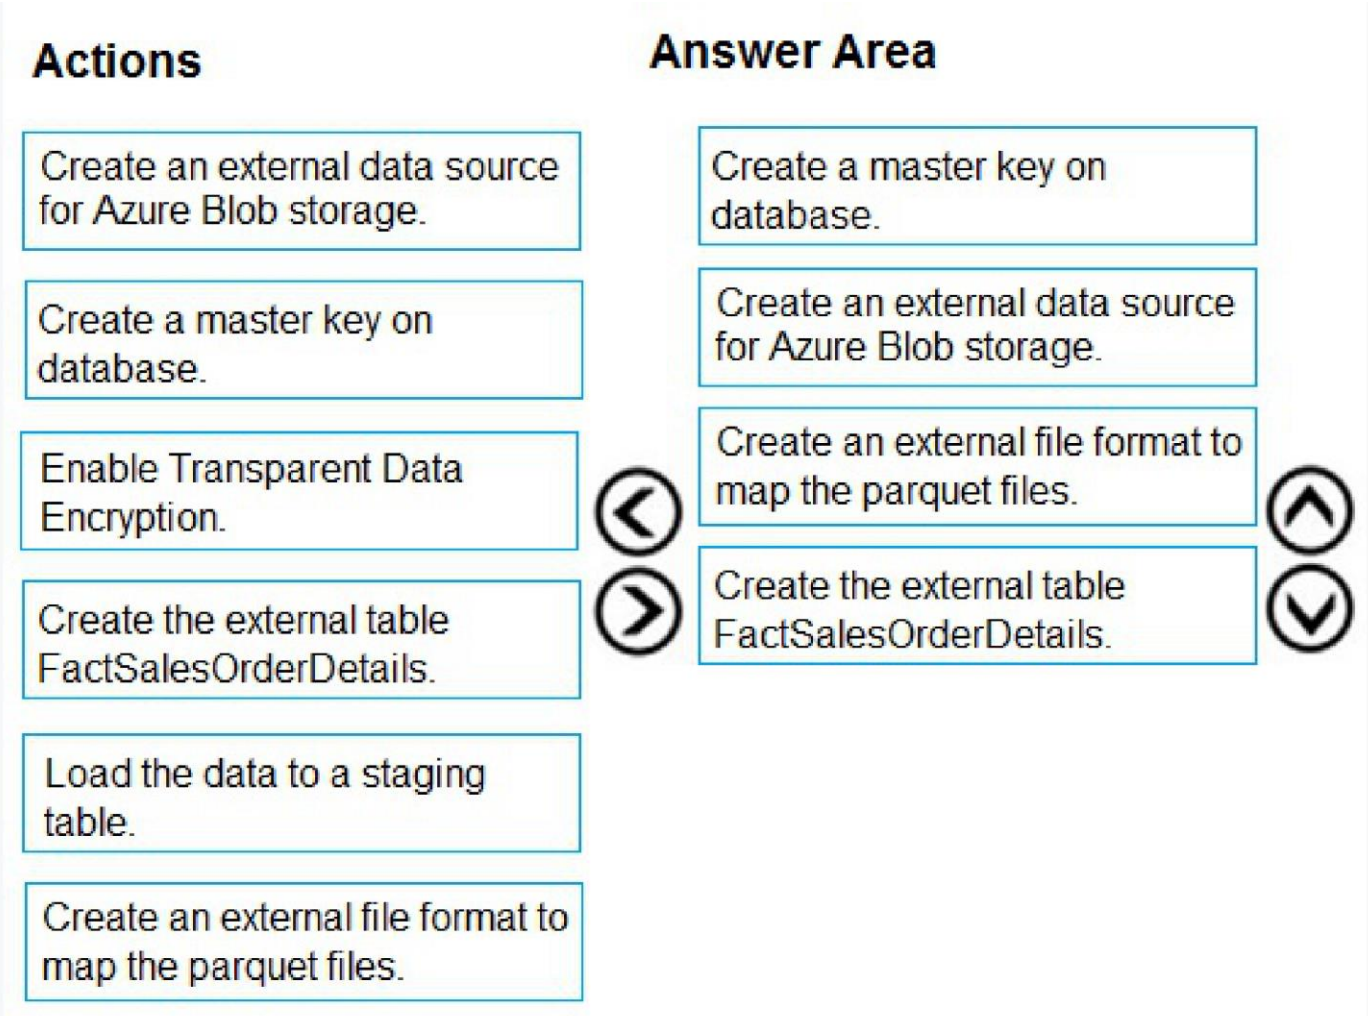

Explanation:
To query the data in your Hadoop data source, you must define an external table to use in Transact-SQL queries.
The following steps describe how to configure the external table.

Step 1: Create a master key on database.

1. Create a master key on the database. The master key is required to encrypt the credential secret.

(Create a database scoped credential for Azure blob storage.) Step 2: Create an external data source for Azure Blob storage.

2. Create an external data source with CREATE EXTERNAL DATA SOURCE..

Step 3: Create an external file format to map the parquet files.

3. Create an external file format with CREATE EXTERNAL FILE FORMAT.

Step 4. Create an external table FactSalesOrderDetails

4. Create an external table pointing to data stored in Azure storage with CREATE EXTERNAL TABLE.

Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-configure-azure-blob-storage

Question: 20

HOTSPOT -
You configure version control for an Azure Data Factory instance as shown in the following exhibit.

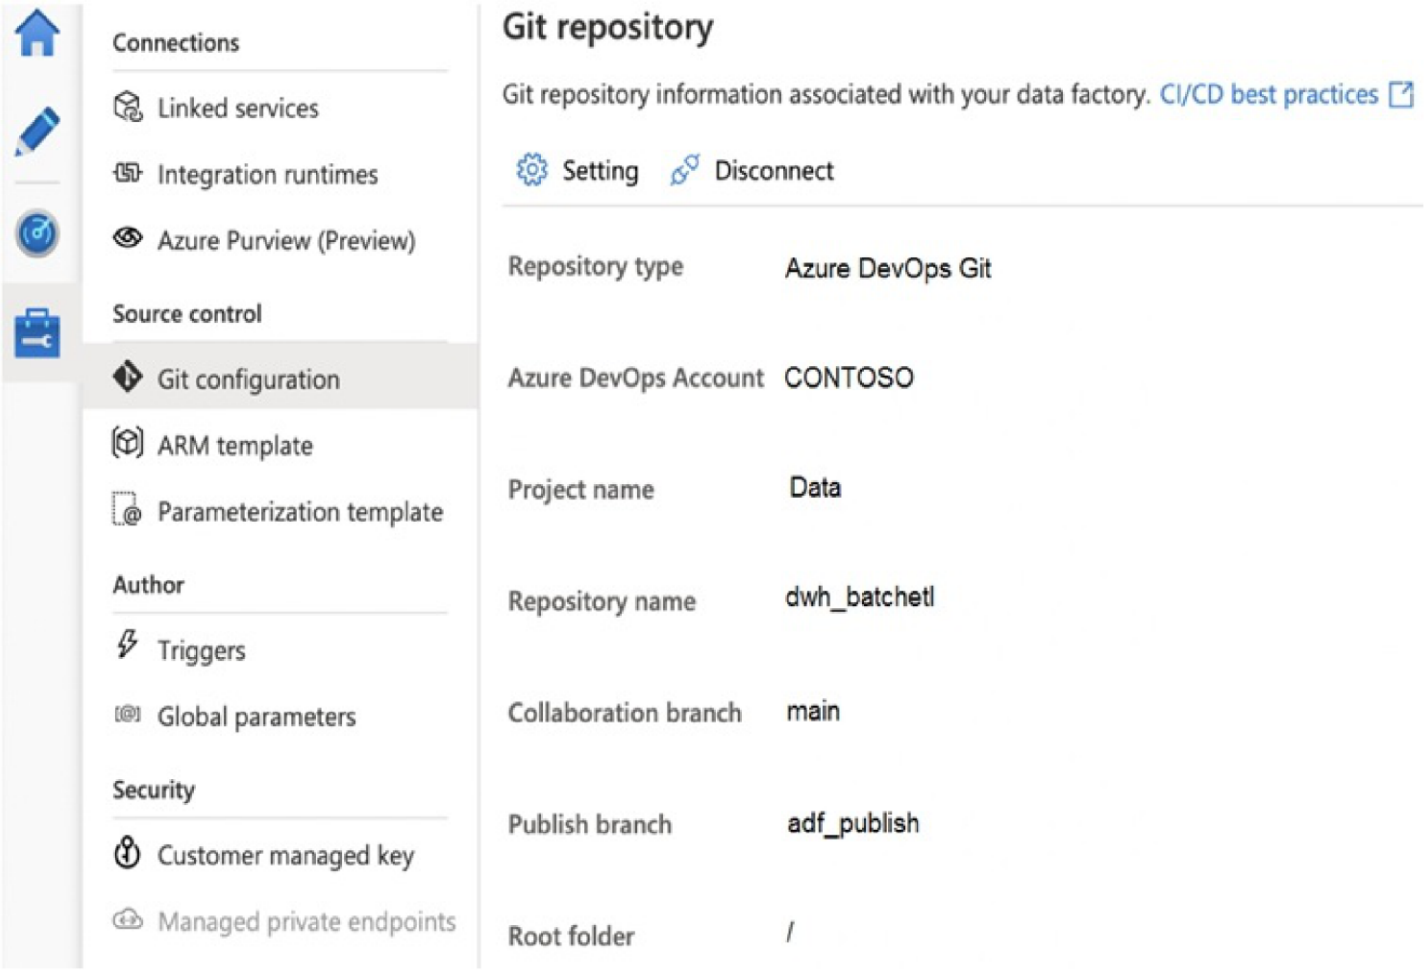

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

Azure Resource Manager (ARM) templates for the pipeline assets as stored in

- /
- adf_publish
- main
- Parameterization template

A Data Factory Azure Resource Manager (ARM) template named contososales can be found in

- /contososales
- /dwh_batchetl/adf_publish/contososales
- /main

Answer:

Answer Area

Azure Resource Manager (ARM) templates for the pipeline assets as stored in

- /
- adf_publish
- main
- Parameterization template

A Data Factory Azure Resource Manager (ARM) template named contososales can be found in

- /contososales
- /dwh_batchetl/adf_publish/contososales
- /main

Explanation:
Box 1: adf_publish -
By default, data factory generates the Resource Manager templates of the published factory and saves them into a branch called adf_publish. To configure a custom publish branch, add a publish_config.json file to the root folder in the collaboration branch. When publishing, ADF reads this file, looks for the field publishBranch, and saves all Resource Manager templates to the specified location. If the branch doesn't exist, data factory will automatically create it. And example of what this file looks like is below:

"publishBranch": "factory/adf_publish"

Box 2: /dwh_barchlet/ adf_publish/contososales
RepositoryName: Your Azure Repos code repository name. Azure Repos projects contain Git repositories to manage your source code as your project grows. You can create a new repository or use an existing repository that's already in your project.

Reference:
https://docs.microsoft.com/en-us/azure/data-factory/source-control

Question: 21

You plan to build a structured streaming solution in Azure Databricks. The solution will count new events in five-minute intervals and report only events that arrive during the interval.
The output will be sent to a Delta Lake table.
Which output mode should you use?

A. complete
B. append
C. update

Answer: B

Explanation:

Append is a correct answer

Question: 22

HOTSPOT -
You are performing exploratory analysis of bus fare data in an Azure Data Lake Storage Gen2 account by using an Azure Synapse Analytics serverless SQL pool.
You execute the Transact-SQL query shown in the following exhibit.

```
SELECT
    payment_type,
    SUM(fare_amount) AS fare_total
FROM OPENROWSET(
            BULK 'csv/busfare/tripdata_2020*.csv',
            DATA_SOURCE = 'BusData',
            FORMAT = 'CSV', PARSER_VERSION = '2.0',
            FIRSTROW = 2
    )
    WITH (
            payment_type INT 10,
            fare_amount FLOAT 11
    ) AS nyc
GROUP BY payment_type
ORDER BY payment_type;
```

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

Answer Area

The query results include only **[answer choice]** in the csv/busfare folder.

- CSV files in the tripdata_2020 subfolder
- files that have files names beginning with "tripdata_2020"
- CSV files that have file names containing "tripdata_202"
- CSV files that have file named beginning with "tripdata_2020"

The query assumes that the first row in a CSV file is **[answer choice]** row.

- a header
- a data
- an empty

Answer:

Answer Area

The query results include only **[answer choice]** in the csv/busfare folder.

- CSV files in the tripdata_2020 subfolder
- files that have files names beginning with "tripdata_2020"
- CSV files that have file names containing "tripdata_202"
- CSV files that have file named beginning with "tripdata_2020" ✓

The query assumes that the first row in a CSV file is **[answer choice]** row.

- a header ✓
- a data
- an empty

Explanation:
Box 1: CSV files that have file named beginning with "tripdata_2020"

Box 2: a header -
FIRSTROW = 'first_row'
Specifies the number of the first row to load. The default is 1 and indicates the first row in the specified data file. The row numbers are determined by counting the row terminators. FIRSTROW is 1-based.

Example: Option firstrow is used to skip the first row in the CSV file that represents header in this case (firstrow=2).
select top 10 * from openrowset(bulk 'https://pandemicdatalake.blob.core.windows.net/public/curated/covid-19/ecdc_cases/latest/ecdc_cases.csv', format = 'csv', parser_version = '2.0', firstrow = 2) as rows

Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-openrowset https://docs.microsoft.co m/en-us/azure/synapse-analytics/sql/query-single-csv-file

Question: 23

You have a SQL pool in Azure Synapse that contains a table named dbo.Customers. The table contains a column name Email.
You need to prevent nonadministrative users from seeing the full email addresses in the Email column. The users must see values in a format of instead.
What should you do?

A. From the Azure portal, set a mask on the Email column.

B. From the Azure portal, set a sensitivity classification of Confidential for the Email column.

C. From Microsoft SQL Server Management Studio, set an email mask on the Email column.

D. From Microsoft SQL Server Management Studio, grant the SELECT permission to the users for all the columns in the dbo.Customers table except Email.

Answer: A

Explanation:

Here ** =** aXXX@XXXX.com

you need to mask the email.

Question: 24

You have an Azure Databricks workspace named workspace1 in the Standard pricing tier. Workspace1 contains an all-purpose cluster named cluster1.
You need to reduce the time it takes for cluster1 to start and scale up. The solution must minimize costs. What should you do first?

A. Upgrade workspace1 to the Premium pricing tier.
B. Configure a global init script for workspace1.
C. Create a pool in workspace1.
D. Create a cluster policy in workspace1.

Answer: C

Explanation:
You can use Databricks Pools to Speed up your Data Pipelines and Scale Clusters Quickly.
Databricks Pools, a managed cache of virtual machine instances that enables clusters to start and scale 4 times faster.

Reference:
https://databricks.com/blog/2019/11/11/databricks-pools-speed-up-data-pipelines.html

Question: 25

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 Synapse Analytics dedicated SQL pool that contains a table named Table1.
You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1. You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.
You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.
Solution: In an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.
Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:
You can use the Get Metadata activity to retrieve the metadata of any data in Azure Data Factory or a Synapse pipeline. You can use the output from the Get
Metadata activity in conditional expressions to perform validation, or consume the metadata in subsequent activities.

Reference:
https://docs.microsoft.com/en-us/azure/data-factory/control-flow-get-metadata-activity

Question: 26

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 Synapse Analytics dedicated SQL pool that contains a table named Table1.
You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1. You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.
You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.
Solution: You use an Azure Synapse Analytics serverless SQL pool to create an external table that has an additional DateTime column.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
This is not about an external table.
Instead, in an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.
Note: You can use the Get Metadata activity to retrieve the metadata of any data in Azure Data Factory or a Synapse pipeline. You can use the output from the
Get Metadata activity in conditional expressions to perform validation, or consume the metadata in subsequent activities.

Reference:
https://docs.microsoft.com/en-us/azure/data-factory/control-flow-get-metadata-activity

Question: 27

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 Synapse Analytics dedicated SQL pool that contains a table named Table1.
You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1. You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.
You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.
Solution: You use a dedicated SQL pool to create an external table that has an additional DateTime column. Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Instead, in an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.
Note: You can use the Get Metadata activity to retrieve the metadata of any data in Azure Data Factory or a Synapse pipeline. You can use the output from the
Get Metadata activity in conditional expressions to perform validation, or consume the metadata in

subsequent activities.

Reference:
https://docs.microsoft.com/en-us/azure/data-factory/control-flow-get-metadata-activity

Question: 28

HOTSPOT -
You are provisioning an Azure SQL database in the Azure portal as shown in the following exhibit.

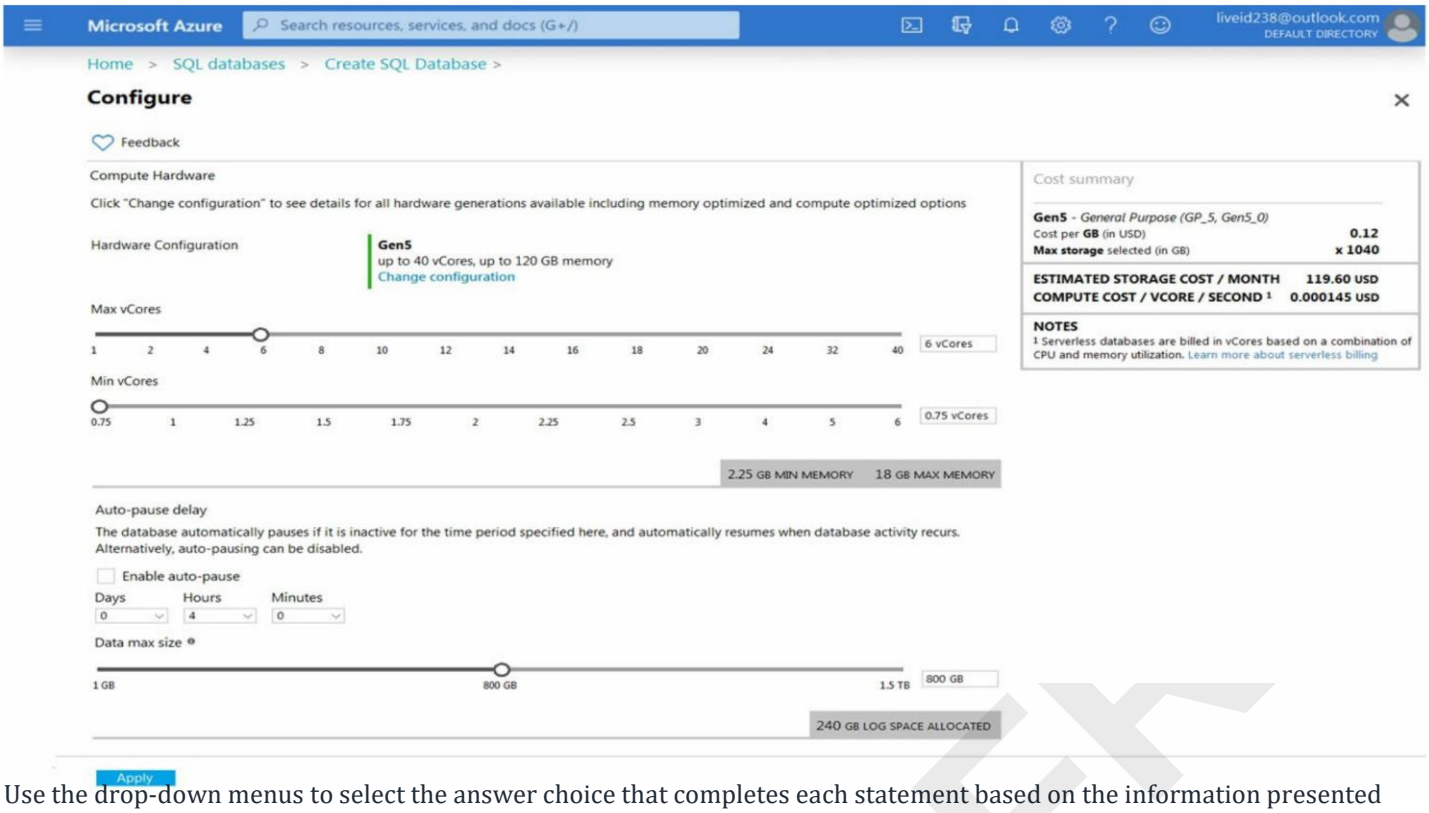

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

After four hours of inactivity, the database requires **[answer choice]** to resume operations for new activities.

- no extra time
- up to 10 minutes
- up to one minute

The database configuration reduces the cost of **[answer choice]** usage patterns.

- intermittent and unpredictable
- regular and high
- steady and low

Answer:

Answer Area

After four hours of inactivity, the database requires **[answer choice]** to resume operations for new activities.

- no extra time
- up to 10 minutes
- up to one minute

The database configuration reduces the cost of **[answer choice]** usage patterns.

- intermittent and unpredictable
- regular and high
- steady and low

Explanation:
Box 1: no extra time -
Auto Pause is not checked in the exhibit.

Note: If Auto Pause is checked the correct answer is: up to one minute Box 2: intermittent and unpredictable

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

Question: 29

You plan to deploy an app that includes an Azure SQL database and an Azure web app. The app has the following requirements:

- The web app must be hosted on an Azure virtual network.
- The Azure SQL database must be assigned a private IP address.
- The Azure SQL database must allow connections only from a specific virtual network. You need to recommend a solution that meets the requirements.

What should you include in the recommendation?

A. Azure Private Link
B. a network security group (NSG)
C. a database-level firewall
D. a server-level firewall

Answer: A

Explanation:
Private Link allows you to connect to various PaaS services in Azure via a private endpoint. A private endpoint is a private IP address within a specific VNet and subnet.

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

Question: 30

You are planning a solution that will use Azure SQL Database. Usage of the solution will peak from October 1 to January 1 each year.
During peak usage, the database will require the following:

- 24 cores

☞ 500 GB of storage
☞ 124 GB of memory
☞ More than 50,000 IOPS
During periods of off-peak usage, the service tier of Azure SQL Database will be set to Standard. Which service tier should you use during peak usage?

A. Business Critical
B. Premium
C. Hyperscale

Answer: A

Explanation:

Business critical in vCore purchasing model) because we have request to scale up and down not only in compute resources but also a storage which is available only vCore purchasing model. Hyperscale is not suits us well as well because it is reversable only within 45 days of the original migration to Hyperscale.

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/database/service-tier-hyperscale?view=azuresql

Question: 31

HOTSPOT -
You have an Azure subscription.
You need to deploy an Azure SQL resource that will support cross database queries by using an Azure Resource Manager (ARM) template.
How should you complete the ARM template? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

```
"resources": [
    ...
    "type":  [ Microsoft.Sql/servers | Microsoft.Sql/servers/databases | Microsoft.Sql/managedInstances ]
     "name": "[parameters('targetName')]",
    "location": "[parameters('location')]",
    "sku": {
     "name": "[parameters('skuName')]"
    ...
    "dependsOn": [
        [ "[parameters('targetName')]" | "[parameters('virtualNetworkName')]" | "[variables('networkSecurityGroupName')]" ],
    ],
    "properties": {
      "administratorLogin": "[parameters('administratorLogin')]",
      "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
      "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets', parameters('virtualNetworkName'),
      parameters('virtualNetworkName'), parameters('subnetName'))]",
     "storageSizeInGB": "[parameters('storageSizeInGB')]","vCores": "[parameters('vCores')]",
      "licenseType": "[parameters('licenseType')]"
...
 ]
```

Answer:

Answer Area

```
"resources": [
    ...
    "type":  [ Microsoft.Sql/servers | Microsoft.Sql/servers/databases | Microsoft.Sql/managedInstances ]
     "name": "[parameters('targetName')]",
    "location": "[parameters('location')]",
    "sku": {
     "name": "[parameters('skuName')]"
    ...
    "dependsOn": [
      [ "[parameters('targetName')]", "[parameters('virtualNetworkName')]" | "[variables('networkSecurityGroupName')]", ]
    ],
    "properties": {
      "administratorLogin": "[parameters('administratorLogin')]",
      "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
      "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets', parameters('virtualNetworkName'),
      parameters('virtualNetworkName'), parameters('subnetName'))]",
     "storageSizeInGB": "[parameters('storageSizeInGB')]","vCores": "[parameters('vCores')]",
      "licenseType": "[parameters('licenseType')]"
...
  ]
```

Explanation:

Box 1: Microsoft.Sql/managedInstances

The Managed Instance depends on the Virtual Network.

Box 2: parameters('virtualNetworkName')

"dependsOn": [

"[resourceId('Microsoft.Network/virtualNetworks', parameters('virtualNetworkName'))]"

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/create-template-quickstart?

tabs=azure-powershell

Question: 32 HOTSPOT

-

You have the following Azure Resource Manager template.

```
...
  "variable": {
    "serverName": "azsqldbserver0001"
  },
  "resources": [
    {
      "name": "[variables('serverName')]",
      "type": "Microsoft.Sql/servers",
      "apiVersion": "2019-06-01-preview",
      "location": "[parameters('location')]",
      "properties": {
        "administratorLogin": "[parameters('administratorLogin')]",
        "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
        "version": "12.0"
      },
      "resources": [
        {
          "name": "[concat(variables('serverName'),'/',parameters('databaseName'))]",
          "type": "Microsoft.Sql/servers/databases",
          "apiVersion": "2020-08-01-preview",
          "location": "[parameters('location')]",
          "kind": "v12.0"
          "sku": {
            "name": "Standard",
            "tier": "Standard",
            "capacity": 10
          },
          "dependsOn": [
           "[concat('Microsoft.Sql/servers/', variables('serverName'))]"
          ],
          "properties": {
          },
          "resources": [
          ]
        }
      ]
    }
  ],
...
```

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:

Answer Area

Statements	Yes	No
The template deploys a serverless Azure SQL database.	○	○
The template deploys a database to an Azure SQL Database managed instance.	○	○
The pricing tier of the database deployment is based on DTUs.	○	○

Answer:

Answer Area

Statements	Yes	No
The template deploys a serverless Azure SQL database.	○	○
The template deploys a database to an Azure SQL Database managed instance.	○	○
The pricing tier of the database deployment is based on DTUs.	○	○

Explanation:

no

no

yes

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/database/purchasing-models
https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-arm-template-quickstart

Question: 33

HOTSPOT -
You have an on-premises Microsoft SQL Server 2019 instance that hosts a database named DB1.
You plan to perform an online migration of DB1 to an Azure SQL managed instance by using the Azure Database Migration Service.
You need to create a backup of DB1 that is accessible to the Azure Database Migration Service.
What should you run for the backup and where should you store the backup? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

Run:
- A full backup and a log backup appended to the same file by using the WITH CHECKSUM option
- A full backup and a log backup to separate files by using the WITH CHECKSUM option
- A full backup and a log backup to separate files by using the WITH FILE_SNAPSHOT option

Store the backup in:
- A Recovery Services vault
- An Azure Blob storage account
- An SMB file share

Answer:

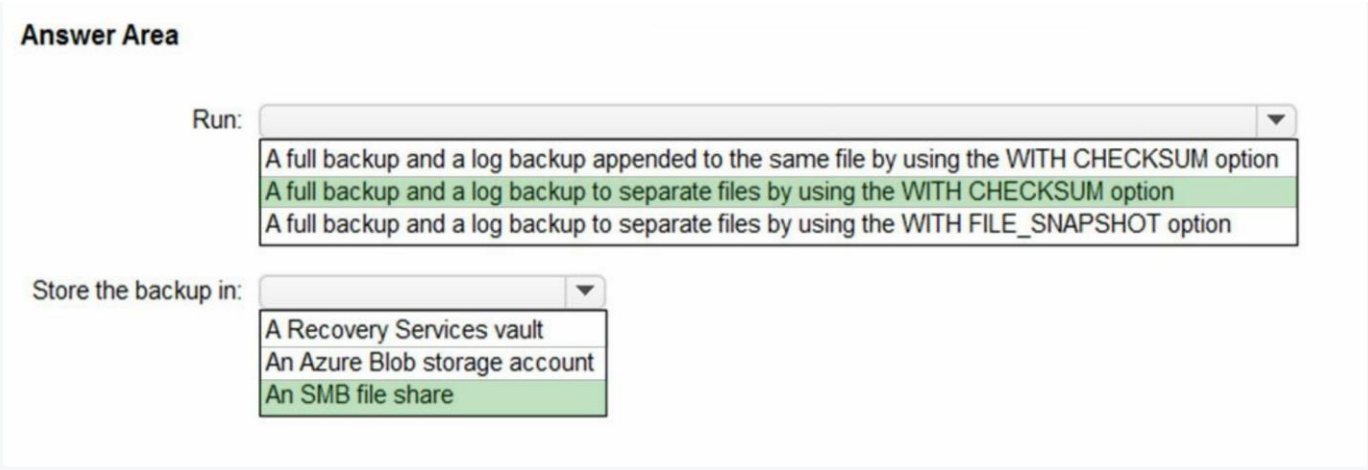

Explanation:
Box 1: ..with CHECKSUM option -
Azure Database Migration Service does not initiate any backups, and instead uses existing backups, which you may already have as part of your disaster recovery plan, for the migration. Be sure that you take backups using the WITH CHECKSUM option.

Box 2: An SMB share -
For online migrations from SQL Server to SQL Managed Instance using Azure Database Migration Service, you must provide the full database backup and subsequent log backups in the SMB network share that the service can use to migrate your databases.

Reference:
https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-managed-instance-online

Question: 34

HOTSPOT -
You have an Azure subscription.
You plan to deploy an Azure SQL database by using an Azure Resource Manager template.
How should you complete the template? To answer, select the appropriate options in the answer area. NOTE:
Each correct selection is worth one point.
Hot Area:

Answer Area

```
{
  "resources": [
    {
      "type": [ "Microsoft.Sql/servers" | "Microsoft.SqlVirtualMachines/sqlVirtualMachines" | "Microsoft.Synapse/workspaces/sqldatabases" ]
      "apiVersion": "2020-02-02-preview",
      "name": "[parameters('name1')]",
      "location": "[parameters('location')]",
    …
      "resources": [
        {
          "type": "databases",
          "apiVersion": "2020-02-02-preview",

    …
          [ "dependsOn":[ | "properties":[ | "tags":[ ]
          "[resourceId('Microsoft.Sql/servers',concat(parameters('name1')))]"
    …
    }
```

Answer:

Answer Area

```
{
  "resources": [
    {
      "type": [ "Microsoft.Sql/servers" | "Microsoft.SqlVirtualMachines/sqlVirtualMachines" | "Microsoft.Synapse/workspaces/sqldatabases" ]
      "apiVersion": "2020-02-02-preview",
      "name": "[parameters('name1')]",
      "location": "[parameters('location')]",
    …
      "resources": [
        {
          "type": "databases",
          "apiVersion": "2020-02-02-preview",

    …
          [ "dependsOn":[ | "properties":[ | "tags":[ ]
          "[resourceId('Microsoft.Sql/servers',concat(parameters('name1')))]"
    …
    }
```

Explanation:
Box 1: "Microsoft.Sql/servers"

Example:
"resources": [

"type": "Microsoft.Sql/servers",
"apiVersion": "2021-08-01-preview",
"name": "[parameters('serverName')]",
"location": "[parameters('location')]",
"properties":
"administratorLogin": "[parameters('administratorLogin')]",
"administratorLoginPassword": "[parameters('administratorLoginPassword')]"

,

"type": "Microsoft.Sql/servers/databases",
"apiVersion": "2021-08-01-preview",
"name": "[format(' 0 / 1 ', parameters('serverName'), parameters('sqlDBName'))]",
"location": "[parameters('location')]",
"sku":
"name": "Standard",
"tier": "Standard"
,
"dependsOn": [
"[resourceId('Microsoft.Sql/servers', parameters('serverName'))]"
]

Box 2: "dependsOn": [

Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-arm-template-quickstart

Question: 35

You have an on-premises Microsoft SQL Server 2019 server that hosts a database named DB1.
You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1 and a virtual network named VNET1. SQLMI1 resides on VNET1.
The on-premises network connects to VNET1 by using an ExpressRoute connection.
You plan to migrate DB1 to SQLMI1 by using Azure Database Migration Service.
You need to configure VNET1 to support the migration.
What should you do?

A. Configure service endpoints.
B. Configure virtual network peering.
C. Deploy an Azure firewall.
D. Configure network security groups (NSGs).

Answer: A

Explanation:
During virtual network setup, if you use ExpressRoute with network peering to Microsoft, add the following service endpoints to the subnet in which the service will be provisioned:
* Target database endpoint (for example, SQL endpoint, Cosmos DB endpoint, and so on)
* Storage endpoint
* Service bus endpoint

This configuration is necessary because Azure Database Migration Service lacks internet connectivity.

Reference:
https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-to-managed-instance

Question: 36

You have an on-premises Microsoft SQL server that uses the FileTables and Filestream features. You plan to migrate to Azure SQL.
Which service should you use?

A. Azure SQL Database
B. SQL Server on an Azure Virtual Machine
C. Azure SQL Managed Instance
D. Azure Database for MySQL

Answer: B

Explanation:
SQL Server VM alternative.

Your business might have requirements that make SQL Server on Azure Virtual Machines a more suitable target than Azure SQL Database.

If one of the following conditions applies to your business, consider moving to a SQL Server virtual machine (VM) instead:
* You have strict dependency on features that are still not supported, such as FileStream/FileTable, PolyBase, and cross-instance transactions.
* Etc.

Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/migration-guides/database/sql-server-to-sql-database-ove rview

Question: 37

You need to migrate an on-premises Microsoft SQL Server database to Azure SQL Database. The solution must minimize downtime.
What should you do?

A. Configure Transaction Log Shipping.
B. Implement Always On availability groups.
C. Configure transactional replication.
D. Import a BACPAC.

Answer: C

Explanation:
Use Transactional Replication.

When you can't afford to remove your SQL Server database from production while the migration is occurring, you can use SQL Server transactional replication as your migration solution.

Note: There are two primary methods for migrating a SQL Server 2005 or later database to Azure SQL Database. The first method (database copy or BACPAC importation) is simpler but requires some, possibly substantial, downtime during the migration. The second method (transactional replication) is more complex,

but substantially eliminates downtime during the migration.

Incorrect:
Not D: The import BACPAC method includes downtime during migration.

Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/migrate-to-database-from-sql-server#method-1-migration-with-downtime-during-the-migration

Question: 38

You have an Azure SQL database named DB1.
You have a table name Table1 that has 20 columns of type CHAR(400). Row compression for Table1 is enabled. During a database audit, you discover that none of the fields contain more than 150 characters.
You need to ensure that you can apply page compression to Table1.
What should you do?

A. Configure the columns as sparse.
B. Change the column type to NVARCHAR(MAX).
C. Change the column type to VARCHAR(MAX).
D. Change the column type to VARCHAR(200).

Answer: D

Explanation:
We reduce the max length of the column from 400 to 200.

Incorrect:
Not A: Sparse column is useful when there are many null columns.

The SQL Server Database Engine uses the SPARSE keyword in a column definition to optimize the storage of values in that column. Therefore, when the column value is NULL for any row in the table, the values require no storage.

Not B, Not C: SQL Server 2005 got around the limitation of 8KB storage size and provided a workaround with varchar(max). It is a non-Unicode large variable- length character data type and can store a maximum of 2^31-1 bytes (2 GB) of non-Unicode characters.

Reference:
https://www.sqlshack.com/sql-varchar-data-type-deep-dive/
https://36chambers.wordpress.com/2020/06/18/nvarchar-everywhere-a-thought-experiment/

Question: 39

You have an on-premises Microsoft SQL Server named SQL1 that hosts five databases.
You need to migrate the databases to an Azure SQL managed instance. The solution must minimize downtime and prevent data loss.
What should you use?

A. Always On availability groups
B. Backup and Restore
C. log shipping
D. Database Migration Assistant

Answer: B

Explanation:

DMA does not support database migrations to Azure SQL Managed Instance. Recommendation is to use the Azure SQL migration extension for Azure Data Studio, which supports both online and offline database migrations to Azure SQL Managed Instance, but here we don´t have the option "Azure SQL migration extension for Azure Data Studio".

Regarding log shipping to Managed Instance is not possible; it only supports the restore of full backups.

https://dba.stackexchange.com/questions/232332/is-it-possible-to-log-ship-from-on-premise-sql-server-to-azure-sql-managed-insta

https://learn.microsoft.com/en-us/sql/dma/dma-overview?view=sql-server-ver16#capabilities

Question: 40

You have an Azure subscription that contains an Azure SQL database. The database contains a table named tablet that uses partitioned columnstores.

You need to configure table1 to meet the following requirements:

• Each partition must be compressed.
• The compression ratio must be maximized.
• You must be able to index the compressed data.

What should you use?

A. page compression
B. columnstore compression
C. GZIP compression
D. columnstore archival compression

Answer: B

Explanation:

While columnstore archival compression can be used to compress columnstore data in a way that minimizes storage costs, it is not optimized for query performance, as it is designed for long-term storage of cold data.

In this scenario, the requirement is to maximize compression ratio while still being able to perform index operations on the compressed data. Columnstore archival compression is not optimized for indexing, and as such, it would not meet this requirement.

Therefore, the correct answer is B. columnstore compression.

Question: 41

You have an Azure subscription linked to an Azure Active Directory (Azure AD) tenant. The subscription contains 10 virtual machines that run Windows Server 2019 and host Microsoft SQL Server 2019 instances.

You need to ensure that you can manage the SQL Server instances by using a single user account.

What should you do first?

A.Enable a user-assigned managed identity on each virtual machine.
B.Deploy an Azure Active Directory Domain Services (Azure AD DS) domain and join the virtual machines to the

domain.

C.Enable a system-assigned managed identity on each virtual machine. D.Join the virtual machines to the Azure AD tenant.

Answer: B

Explanation:

Deploy an Azure Active Directory Domain Services (Azure AD DS) domain and join the virtual machines to the domain.

Question: 42 DRAG DROP

-

You have an Azure subscription.

You plan to deploy a new Azure virtual machine that will host a Microsoft SQL Server instance.

You need to configure the disks on the virtual machine. The solution must meet the following requirements:• Minimize latency for transaction logs.

• Minimize the impact on IO throughput of the virtual machine.

Which type of disk should you use for each workload? To answer, drag the appropriate disk types to the correct workloads. Each disk type 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.

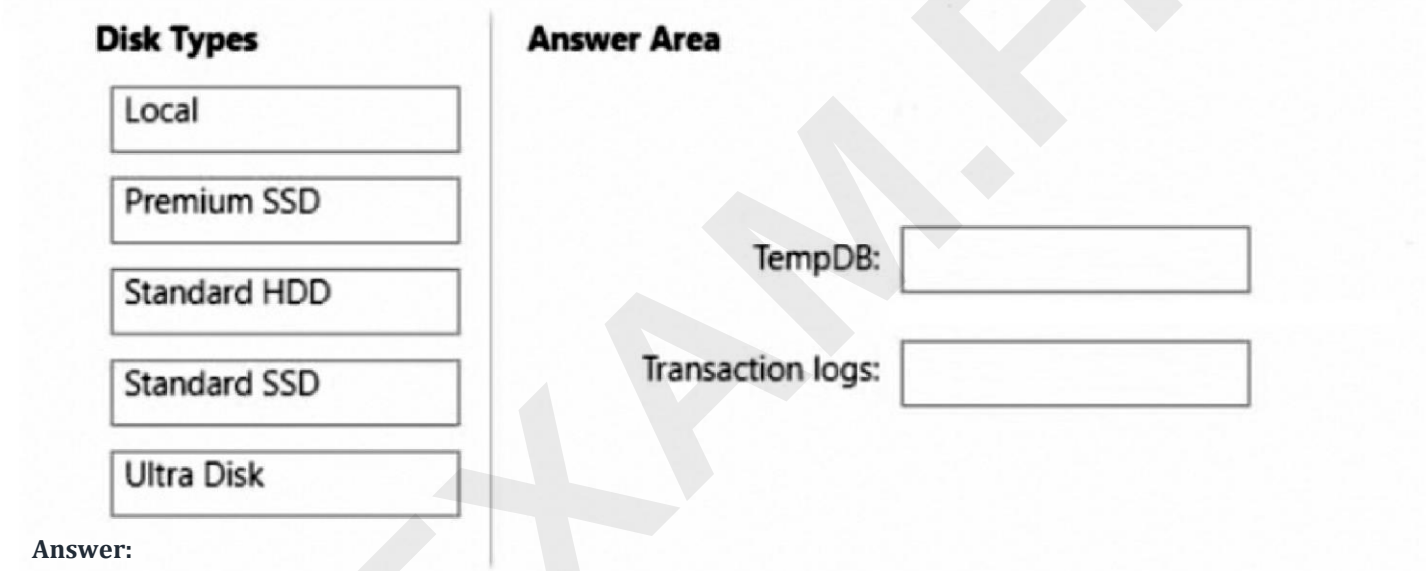

Answer:

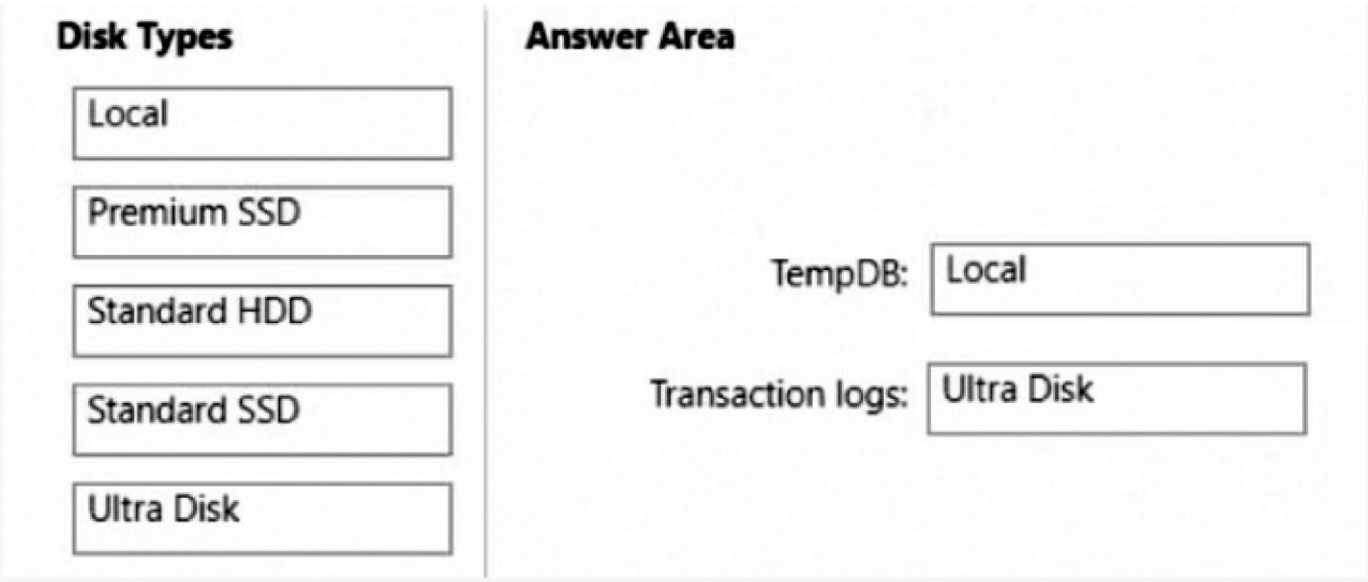

Explanation:

Temp DB: Local.

Transaction logs: Ultra Disk.

Question: 43

You have an Azure SQL Database elastic pool that contains 10 databases.

You receive the following alert.

Msg 1132, Level 16, state 1, Line 1
The elastic pool has reached its storage limit. The storage used for the elastic pool cannot exceed (76800) MBs. You need to resolve the alert. The solution must minimize administrative effort.

Which three actions can you perform? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A.Increase the maximum storage of the elastic pool.
B.Delete data from a database.
C.Remove a database from the pool.
D.Enable data compression.
E.Shrink individual databases.

Answer: ACE

Explanation:

A.Increase the maximum storage of the elastic pool.

C.Remove a database from the pool.

E.Shrink individual databases.

Question: 44

You have an Azure subscription.

You need to deploy a new Azure SQL database by using Azure Command-Line Interface (CLI). Which three parameters are required?

A.--name, --edition, and --capacity
B.--name, --tier, and --min-capacity
C.--name, --resource-group, and --server
D.--name, --licence-type, and --capacity

Answer: C

Explanation:

--name, --resource-group, and --server

Question: 45 HOTSPOT

-

You have an Azure subscription.

You plan to migrate 10 on-premises Microsoft SQL Server instances to Azure.

You need to ensure that the migrated environment can be managed by using multiserver administration and supports master/target (MSX/TSX) jobs.

The solution must minimize administrative effort.

Which SQL deployment options should you select as the master server (MSX) and the target server (TSX)? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

MSX:

- SQL database
- SQL managed instances
- SQL virtual machines

TSX:

- SQL database
- SQL managed instances
- SQL virtual machines

Answer:

Answer Area

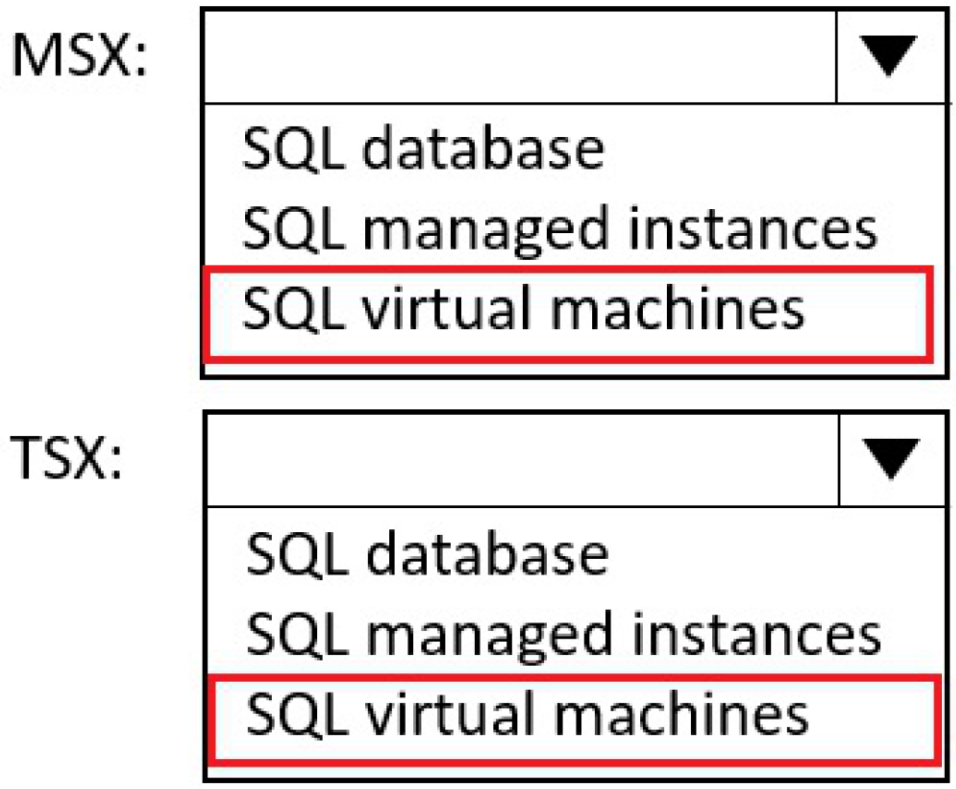

Explanation:
MSX : SQL Virtual Machines.

SQL Virtual Machines offer full control and compatibility for managing multiserver jobs via MSX/TSX. TSX : SQL Virtual Machines.

Only SQL Virtual Machines can function as a target for multiserver jobs in Azure.

Question: 46 You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2.

You need to migrate the databases hosted on SQL1 to Azure. The solution must meet the following requirements:• The service that hosts the migrated databases must be able to communicate with SQL2 by using linked server connections.
• Administrative effort must be minimized.

What should you use to host the databases?

A. a single Azure SQL database
B. SQL Server on Azure Virtual Machines
C. Azure SQL Managed Instance
D. an Azure SQL Database elastic pool

Answer: C

Explanation:

Azure SQL Managed Instance

Question: 47 HOTSPOT

-

You have an on-premises Microsoft SQL Server 2016 instance that hosts a database named db1. You have an Azure subscription that contains an Azure SQL managed instance named MI1.

You plan to perform an online migration of db1 to MI1 by using Azure Database Migration Service.

You need to create the backups for the migration. The solution must minimize the number of backup files created.

Which type of backups should you create, and how should you store the backups? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Backup type:

- Full only
- Full and differential
- Full and transaction log
- Transaction log only

To store the backups:

- Append all the backups to a single file.
- Append each backup to a separate file.
- Write each backup to a separate file.

Answer:

Answer Area

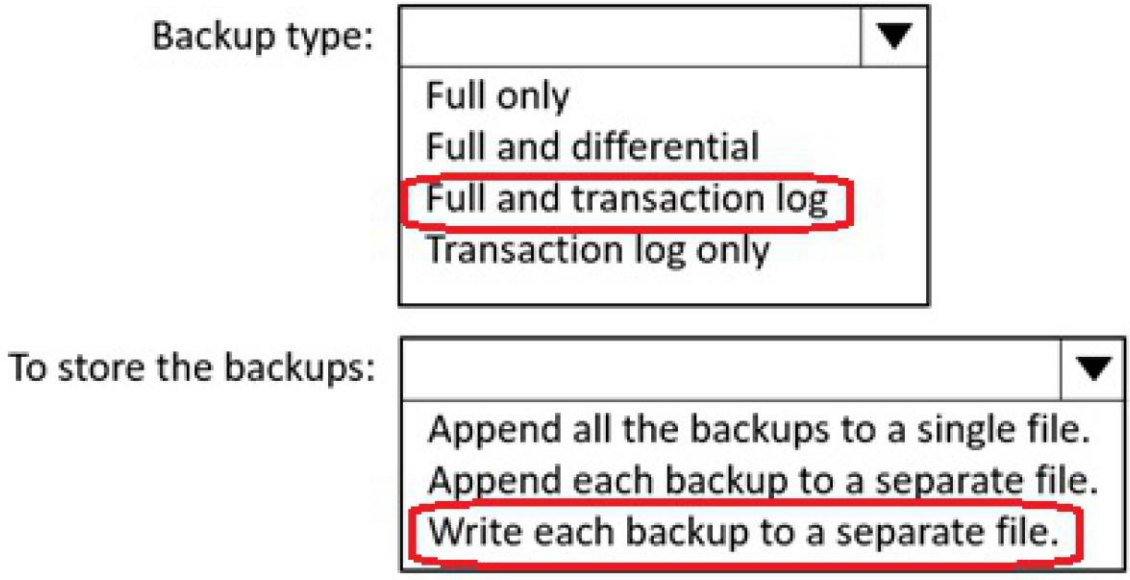

Question: 48 You have a SQL Server on Azure Virtual Machines instance named SQLVM1 that was deployed by using an Azure Marketplace SQL Server 2019 Enterprise image.

You need to change the Microsoft SQL Server instance on SQLVM1 to the Standard edition. The solution must ensure licensing compliance.

What should you do first?

A. From the SQL Server Installation Center on SQLVM1, run the Edition Upgrade wizard. B. From SQLVM1, uninstall the SQL Server instance.

C. From the SQL Server Installation Center on SQLVM1, run the Repair wizard. D. From the Azure portal, reconfigure SQLVM1.

Answer: B

Explanation:

From SQLVM1, uninstall the SQL Server instance.

Question: 49 Your on-premises network contains a Microsoft SQL Server 2016 server that hosts a database named db1. You have an Azure subscription.

You plan to migrate db1 to an Azure SQL managed instance.

You need to create the SQL managed instance. The solution must minimize the disk latency of the instance. Which service tier should you use?

A. Business Critical

B. Hyperscale
C. General Purpose
D. Premium

Answer: A

Explanation:

Business Critical

Question: 50 You have an Azure subscription.

You need to deploy an Azure SQL database. The solution must meet the following requirements:• Dynamically scale CPU resources.
• Ensure that the database can be paused to reduce costs.

What should you use?

A. the Business Critical service tier
B. the serverless compute tier
C. an elastic pool
D. the General Purpose service tier

Answer: B

Explanation:

the serverless compute tier

Question: 51 HOTSPOT -

You have an Azure subscription.

You need to deploy an Azure SQL managed instance that meets the following requirements:•Optimize latency.
•Maximize the memory-to-vCore ratio.

Which service tier and hardware generation should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

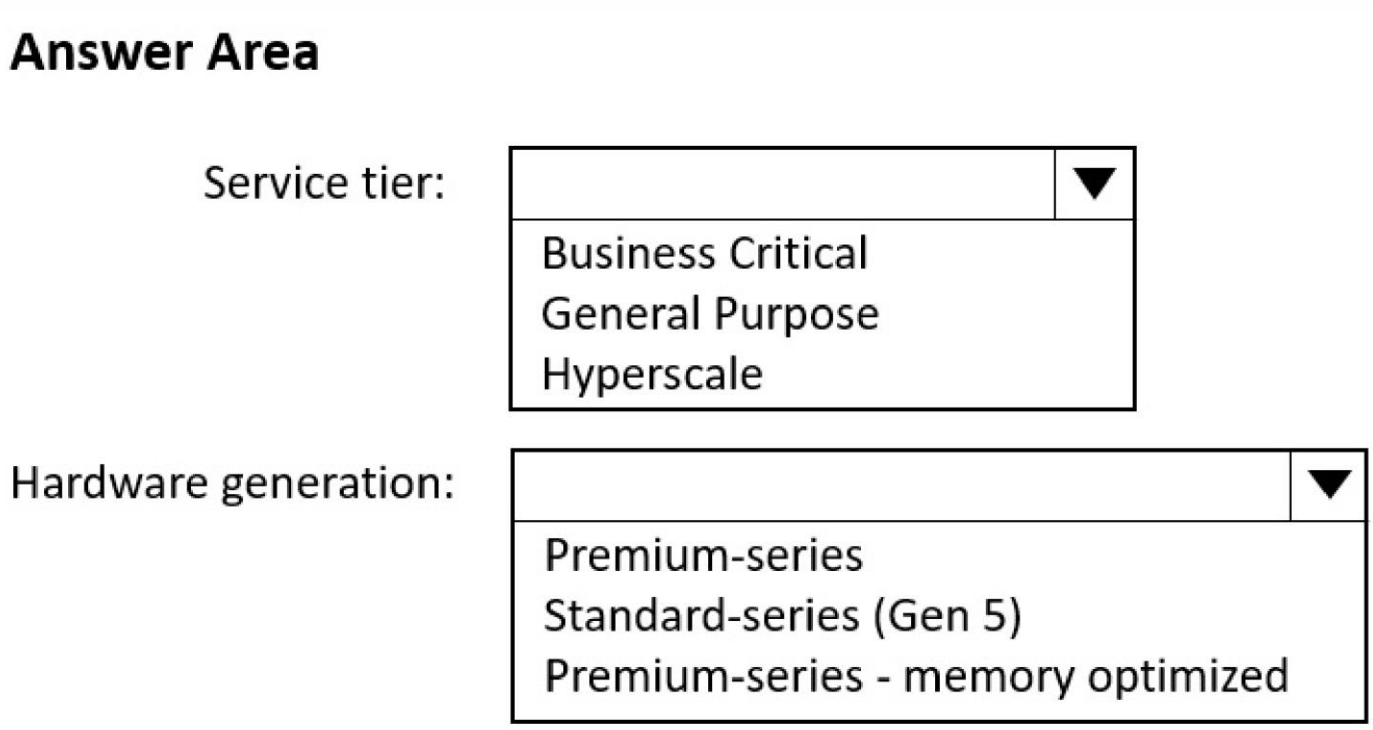

Answer:

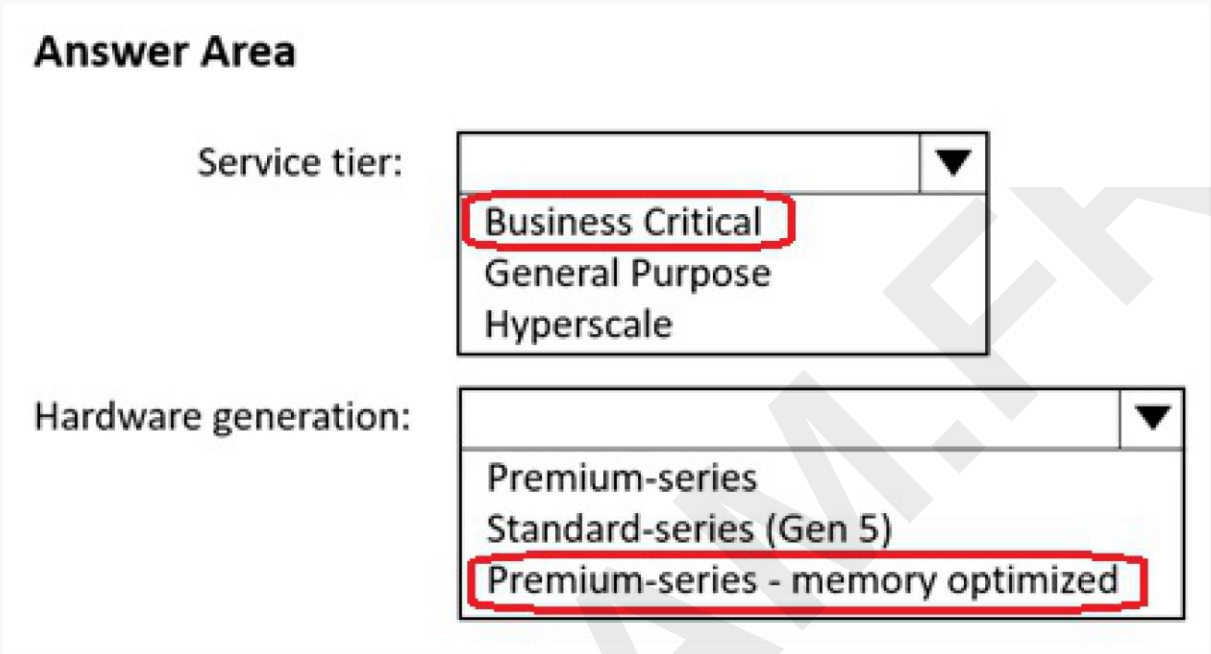

Question: 52

You have a Microsoft SQL Server 2017 server.

You need to migrate the server to Azure. The solution must meet the following requirements:•Ensure that the latest version of SQL Server is used.
•Support the SQL Server Agent service.
•Minimize administrative effort.

What should you use?

A.an Azure SQL Database elastic pool
B.Azure SQL Database

C.SQL Server on Azure Virtual Machines
D.Azure SQL Managed Instance

Answer: D

Explanation:

The answer is D. Because: ensure that the latest version of SQL Server is used.Only SQL MI is always in latest version. SQL on VM can be in previus version.

Question: 53 HOTSPOT

-

You have a Microsoft SQL Server 2017 server that hosts five databases.

You plan to migrate the databases to Azure.

You need to recommend a solution that meets the following requirements:

•Automatically scales compute based on the workload demand

•Provides per-second billing

What should you include in the recommendation? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

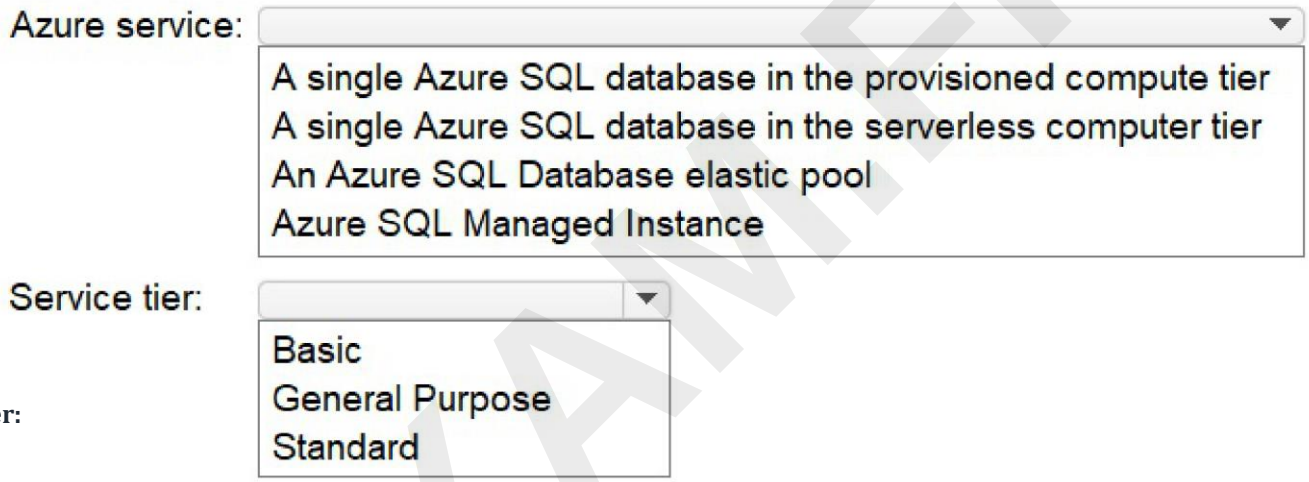

Answer:

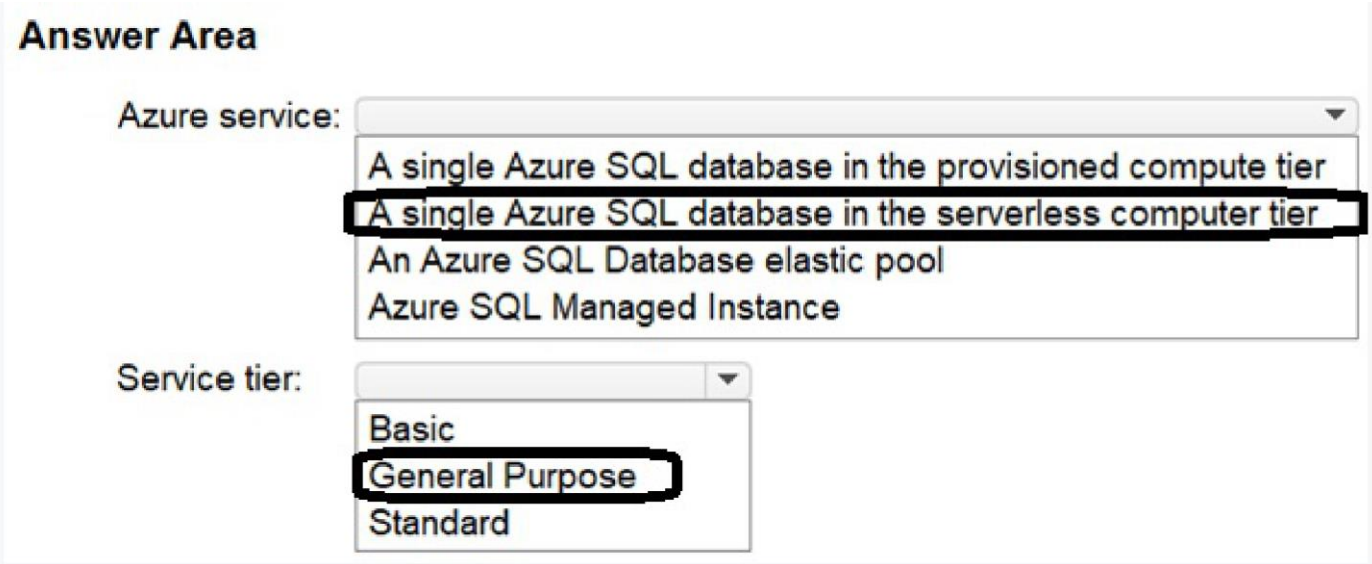

Question: 54

You have an on-premises Microsoft SQL Server 2019 database named SQL1 that uses merge replication.

You need to migrate SQL1 to Azure.

Which service should you use?

A.Azure SQL Edge
B.Azure SQL Database
C.SQL Server on Azure Virtual Machines
D.Azure SQL Managed Instance

Answer: C

Explanation:

SQL Server on Azure Virtual Machines

Question: 55

You have an on-premises datacenter that contains a 2-TB Microsoft SQL Server 2019 database named DB1.

You need to recommend a solution to migrate DB1 to an Azure SQL managed instance. The solution must minimize downtime and administrative effort.

What should you include in the recommendation?

A.Log Replay Service (LRS)
B.log shipping
C.transactional replication
D.SQL Data Sync

Answer: A

Explanation:

Answer A. Less administrative effort compared to transactional replication

https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/log-replay-service-migrate?view=azuresql&tabs=sas-token

Question: 56 You have an Azure subscription.

You plan to deploy an instance of SQL Server on Azure Virtual Machines that supports Write Accelerator. Which virtual machine series should you use?

A.E-series
B.G-series
C.H-series
D.M-series

Answer: D

Explanation:

Answer D

https://learn.microsoft.com/en-us/azure/virtual-machines/how-to-enable-write-accelerator

Question: 57 You have an on-premises Microsoft SQL Server 2019 instance that hosts a database named DB1.
You have an Azure subscription that contains an Azure SQL database named SQLDB1.

You need to replicate DB1 to SQLDB1.

Which type of replication should you use?

A.transactional
B.peer-to-peer
C.snapshot
D.merge

Answer: A

Explanation:

transactional is a correct answer.

Question: 58 HOTSPOT -

You have two on-premises servers that run Windows Server 2019 and host a Microsoft SQL server 2017 Always On availability group named AG1. AG1 contains a single database named DB1.

You have an Azure subscription. The subscription contains a virtual machine named VM1 that runs Linux. You need to migrate DB1 to a SQL Server 2019 instance on VM1. The solution must minimize the downtime of DB1 during the migration.

What should you do? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one-point.

Answer Area

To prepare for the migration:

- Add a secondary replica to AG1.
- Create a SQL Server 2019 Always On availability group on VM1.
- Upgrade the on-premises SQL servers to SQL Server 2019.

To perform the migration, use:

- A distributed availability group
- Azure Migrate
- Log shipping

Answer:

Answer Area

To prepare for the migration: Add a secondary replica to AG1.

To perform the migration, use: Azure Migrate

Question: 59

DRAG DROP

-

You have an Azure SQL database named DB1.

You need to create a partitioned table in DB1.

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

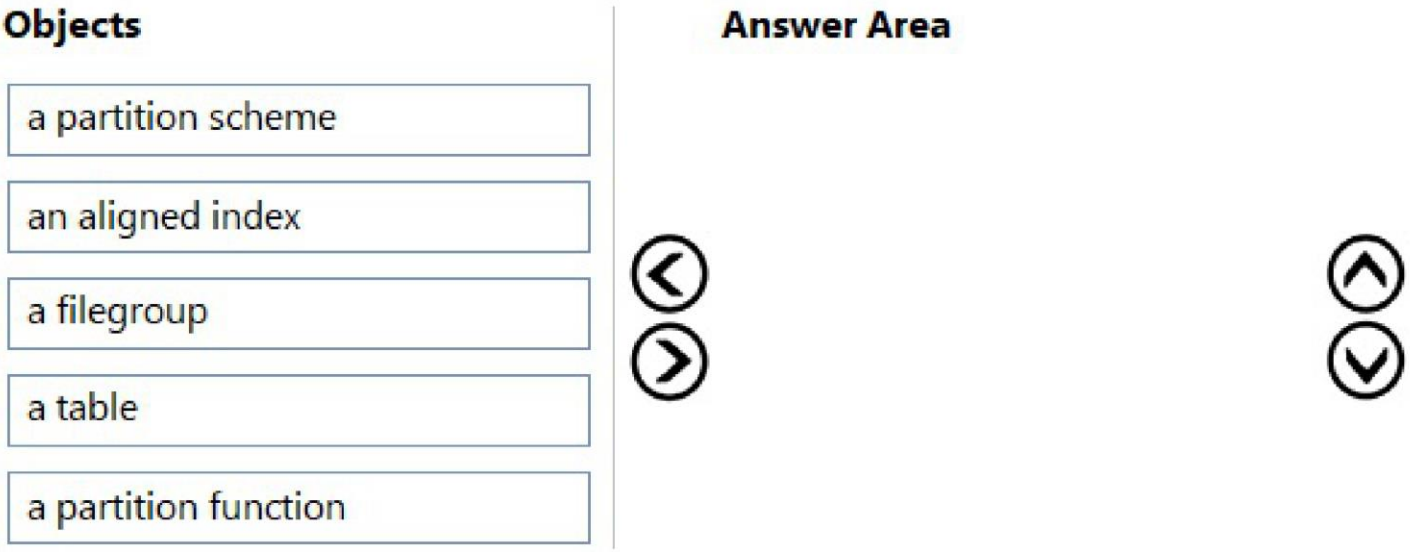

Answer:

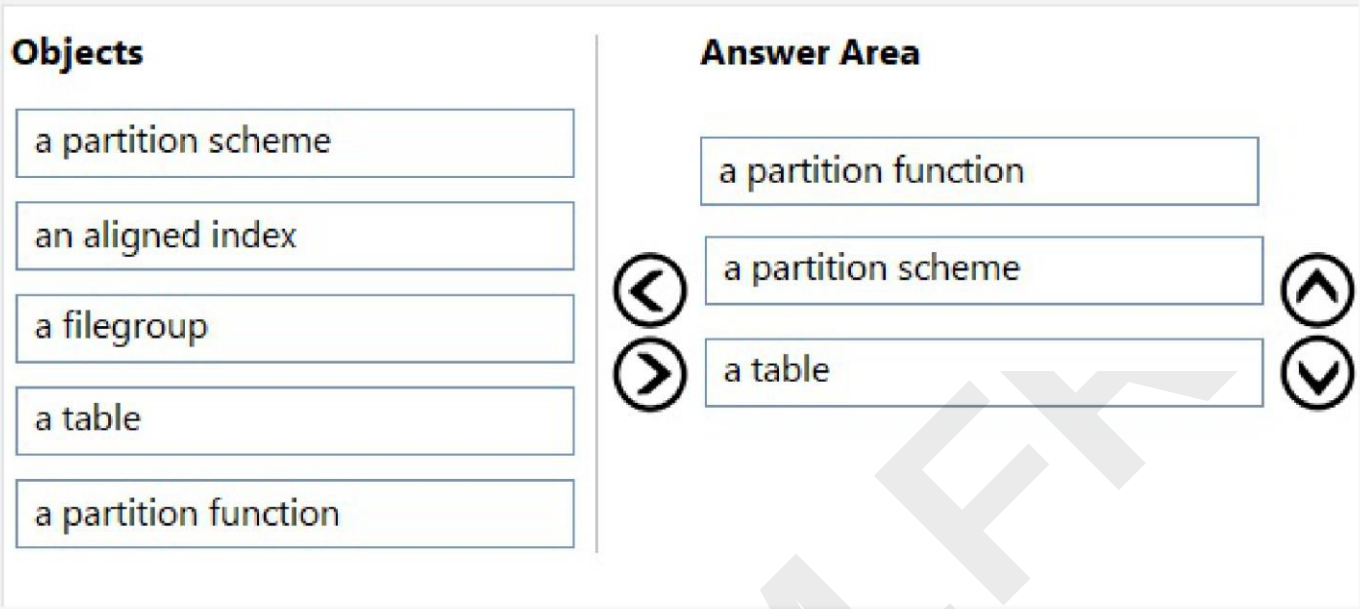

Explanation:

A partition functionA partition schemeA table

Question: 60

DRAG DROP
-

You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines named SQLVM1 and a virtual machine named Server1 that runs Windows Server. SQLVM1 and Server1 are joined to an Active Directory Domain Services (AD DS) domain. Server1 hosts a file share named Share1.

You need to ensure that a SOL Server Agent job step on SQLVM1 can access the files in Share1. The solution must use the principle of least privilege.

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.

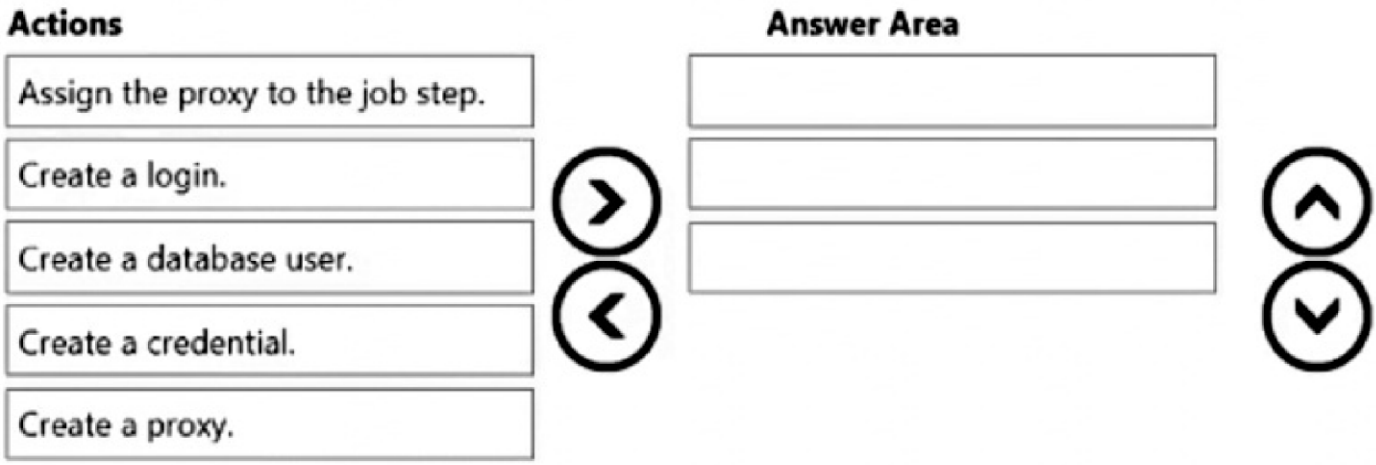

Answer:

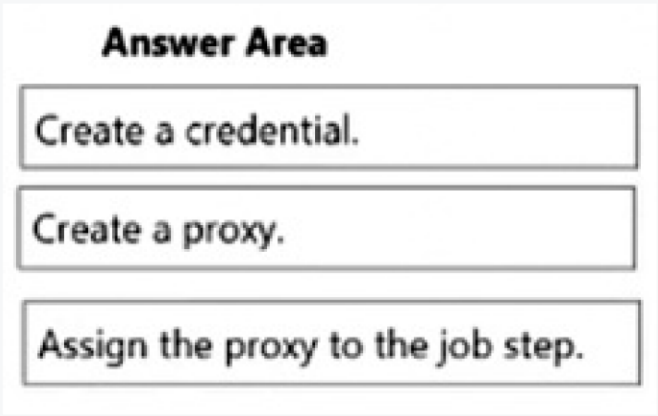

Question: 61

You have an Azure subscription.

You need to deploy an instance of SQL Server on Azure Virtual Machines. The solution must meet the following requirements:
•Custom performance configuration, such as IOPS, capacity, and throughout, must be supported.
•Costs must be minimized.

Which type of disk should you include in the solution?

A.Premium SSD v2
B.Premium SSD
C.Standard SSD
D.Ultra SSD

Answer: D

Explanation:

Correct answer is D:Ultra SSD.

Question: 62

You have an on-premises datacenter that contains a 14-TB Microsoft SQL Server database.

You plan to create an Azure SQL managed instance and migrate the on-premises database to the new instance. Which three service tiers support the SQL managed instance? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A.General Purpose Standard
B.Business Critical Memory Optimized Premium
C.General Purpose Premium
D.Business Critical Premium
E.Business Critical Standard

Answer: ABC

Explanation:

A.General Purpose Standard

B.Business Critical Memory Optimized Premium

C.General Purpose Premium

Question: 63

SIMULATION

-

You need to configure db1 to pause automatically after one hour of inactivity.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

Azure SQL elastic pool pausing feature
Aan elastic pool vCore provisioned server can be paused. You can do this by changing the compute tier of the elastic pool to Serverless. In the serverless compute tier, the database is automatically paused when it is inactive for a period of time that you specify. The database is automatically resumed when the next login or other activity occurs.

Part 1: Change to the serverless compute tier

To change the compute tier of an elastic pool, you can follow these steps:
Step 1: Go to the Azure portal and sign in to your account.

Step 2: Select the Azure SQL resource that contains the elastic pool.

Step 3: In the left navigation menu, select Elastic Pools.

Step 4: Select the elastic pool that you want to change the compute tier for.

Step 5: In the Settings section, select Compute Tier.

Step 6: Select Serverless.

Step 7: Click Save.

Part 2: Configure the Auto-pause setting.

Step 1: Select db1.
Once you have changed the compute tier to serverless, the database will be paused after the idle time that you specified. You can view the idle time in the Auto-pause delay setting.

Step 2: Open database settings.

Step 3: Change the Auto-pause delay setting to 1 hour.

Reference:
https://learn.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview
https://learn.microsoft.com/en-us/azure/azure-sql/database/single-database-create-quickstart

Question: 64

SIMULATION
-

You need to ensure that any enhancements made to the Query Optimizer through patches are available to db1 and db2 on sql12345678.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

Stored procedure execution
You can use the SQL Server Management Studio (SSMS) user interface or Transact-SQL in an SSMS query window to execute a stored procedure. Always use the latest version of SSMS.

Use SQL Server Management Studio

Step 1: In Object Explorer, connect to an instance of SQL Server or Azure SQL Database [here sql12345678], expand that instance, and then expand Databases.

Step 2: Expand the database that you want [Here first db1], expand Programmability, and then expand Stored Procedures.

Step 3: Right-click the stored procedure that you want to run [Here: ALTER DATABASE SCOPED CONFIGURATION] and select Execute Stored Procedure.

Step 4: Add parameter: QUERY_OPTIMIZER_HOTFIXES = ON

Step 5: Select OK to execute the stored procedure. If the stored procedure doesn't have any parameters, just select OK.

Step 6: Repeated step 2 to step 5 for db2.

Note: The default approach that SQL Server uses for the query optimizer enables any fixes for the latest database compatibility level for a given product release. This means any fixes for the query optimizer will be used up to the compatibility level of the database, but any hotfixes beyond that compatibility level will not be used. Microsoft provides an option to enable query optimizer hotfixes using the scoped configuration QUERY_OPTIMIZER_HOTFIXES.

Use the following command to enable optimizer hotfixes:

ALTER DATABASE SCOPED CONFIGURATION QUERY_OPTIMIZER_HOTFIXES = ON;

This command configures the database to use all query optimizer hotfixes. For example, optimizations that were included in a recent database cumulative update will apply if the administrator has altered the scope to include query optimizer hotfixes. Using QUERY_OPTIMIZER_HOTFIXES applies at the database level, meaning the change would have to be made for each database.

Reference:
https://infohub.delltechnologies.com/en-US/l/performance-best-practices-4/sql-server-2019-query-optimizer-hotfixes
https://learn.microsoft.com/en-us/sql/t-sql/statements/alter-database-scoped-configuration-transact-sql

Question: 65

SIMULATION
-

You need to add an Azure AD user named to db1. User2-12345678 must be able to read data from all the tables in db1 without being able to modify the data.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

Stage 1: Create the user

Create a user with SSMS (SQL Server Management Studio)

Step 1: In Object Explorer, expand the Databases folder.

Step 2: Expand the database in which to create the new database user. [Here db1]

Step 3: Right-click the Security folder, point to New, and select User....

Step 4: In the Database User - New dialog box, on the General page, select one of the following user types from the User type list:

SQL user with login
SQL user with password (when contained database is enabled)
SQL user without login
User mapped to a certificate
User mapped to an asymmetric key
*-> Windows user

Step 4a: Select Windows User

Step 5: When you select an option, the remaining options in the dialog may change. Some options only apply to specific types of database users. Some options can be left blank and will use a default value.

User name
Enter a name for the new user. If you have chosen Windows user from the User type list, you can also select the ellipsis (...) to open the Select User or Group dialog box. [Specify User2-12345678]

Step 6: Select OK.

Stage 2: Assign permissions

Additional Options
The Database User - New dialog box also offers options on four other pages: Owned Schemas, Membership, Securables, and Extended Properties.

Step 7: Select Securables

The Securables page lists all possible securables and the permissions on those securables that can be granted to the login.

Step 8: Select tables

Step 9: Grant Read Permission

Reference:
https://learn.microsoft.com/en-us/sql/relational-databases/security/authentication-access/create-a-database-user

Question: 66

HOTSPOT
-

Case study
-

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study
-
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview
-
ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment. Licensing Agreement
ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Existing Environment. Network Infrastructure
ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Existing Environment. Identity Environment
The on-premises network contains an Active Directory Domain Services (AD DS) forest. The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Existing Environment. Database Environment
The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1. App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements. Planned Changes
-

ADatum plans to implement the following changes:

•Deploy an Azure SQL managed instance named Instance1 to Network1.
•Migrate DB1 and DB2 to Instance1.
•Migrate DB3 to Azure SQL Database.
•Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
•Following the migration of DB3, configure the database to be part of an auto-failover group.

Requirements. Availability Requirements

ADatum identifies the following post-migration availability requirements:

•For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.•Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
•After the migration, App1 must maintain access to DB1 and DB2.
•For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.•Ensure that DB3 will still be accessible following a planned failover.
•Ensure that DB3 can be restored if the logical server is deleted.
•Minimize downtime during the migration of DB1 and DB2.

Requirements. Security Requirements

ADatum identifies the following security requirements for after the migration:

•Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
•Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Requirements. Management Requirements

ADatum identifies the following post-migration management requirements:

•Continue using Extended Events to monitor DB3.
•In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

Requirements. Business Requirements

ADatum identifies the following business requirements:

•Minimize costs whenever possible, without affecting other requirements.•Minimize administrative effort.

You need to recommend which service and target endpoint to use when migrating the databases from SVR1 to Instance1. The solution must meet the availability requirements.

What should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

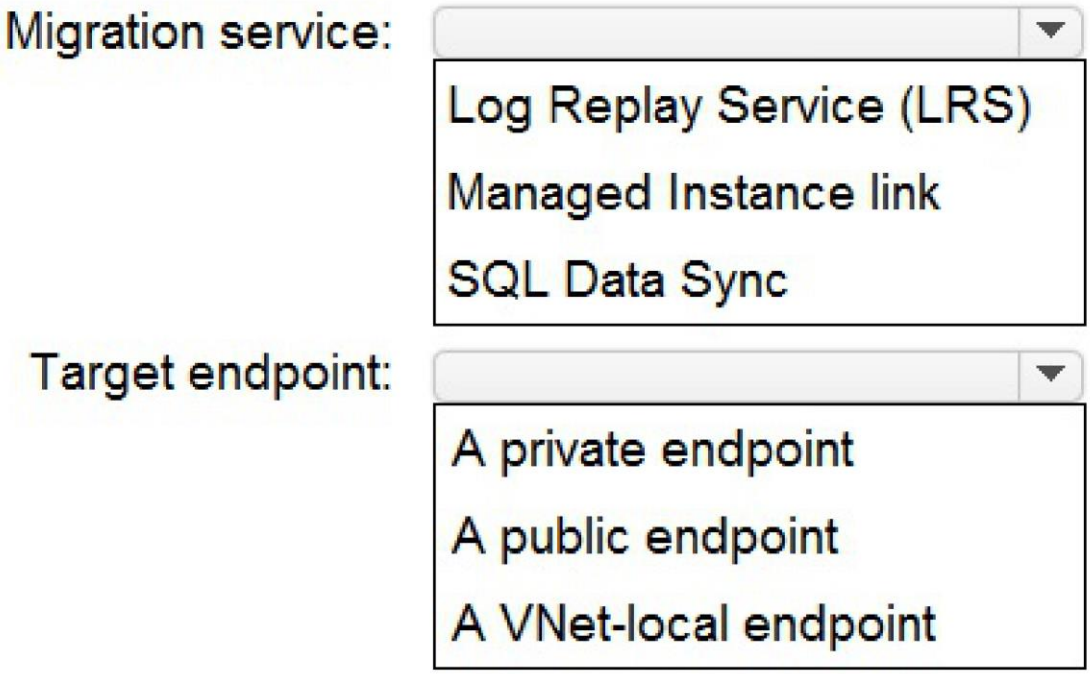

Answer:

Answer Area

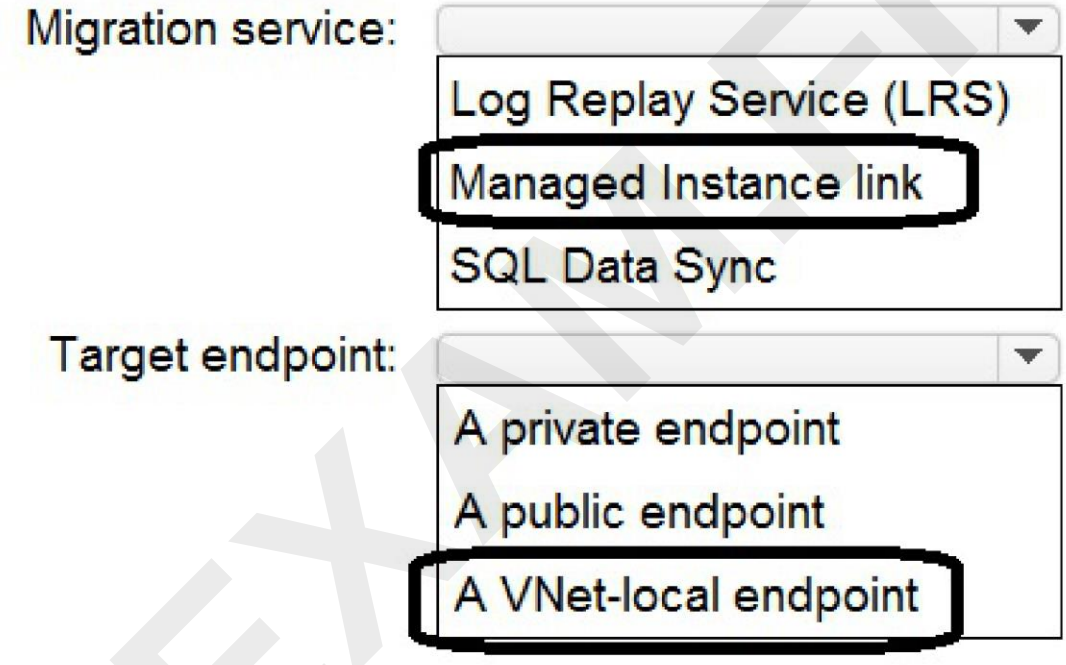

Explanation:

Managed instance link.

A VNet-local endpoint.

Question: 67 HOTSPOT

-

Case study

-

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

-

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

-

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment. Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Existing Environment. Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Existing Environment. Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest. The

forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Existing Environment. Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1.

App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements. Planned Changes
-

ADatum plans to implement the following changes:

•Deploy an Azure SQL managed instance named Instance1 to Network1.
•Migrate DB1 and DB2 to Instance1.
•Migrate DB3 to Azure SQL Database.
•Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
•Following the migration of DB3, configure the database to be part of an auto-failover group.

Requirements. Availability Requirements

ADatum identifies the following post-migration availability requirements:

•For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.•Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
•After the migration, App1 must maintain access to DB1 and DB2.
•For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.•Ensure that DB3 will still be accessible following a planned failover.
•Ensure that DB3 can be restored if the logical server is deleted.
•Minimize downtime during the migration of DB1 and DB2.

Requirements. Security Requirements

ADatum identifies the following security requirements for after the migration:

•Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
•Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Requirements. Management Requirements

ADatum identifies the following post-migration management requirements:

•Continue using Extended Events to monitor DB3.
•In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped

credentials.

Requirements. Business Requirements
ADatum identifies the following business requirements:
•Minimize costs whenever possible, without affecting other requirements.•Minimize administrative effort.

You need to recommend a service tier and a method to offload analytical workloads for the databases migrated from SVR1. The solution must meet the availability and business requirements.

What should you recommend? To answer, select the appropriate options in the answer area. NOTE:
Each correct selection is worth one point.

Answer Area

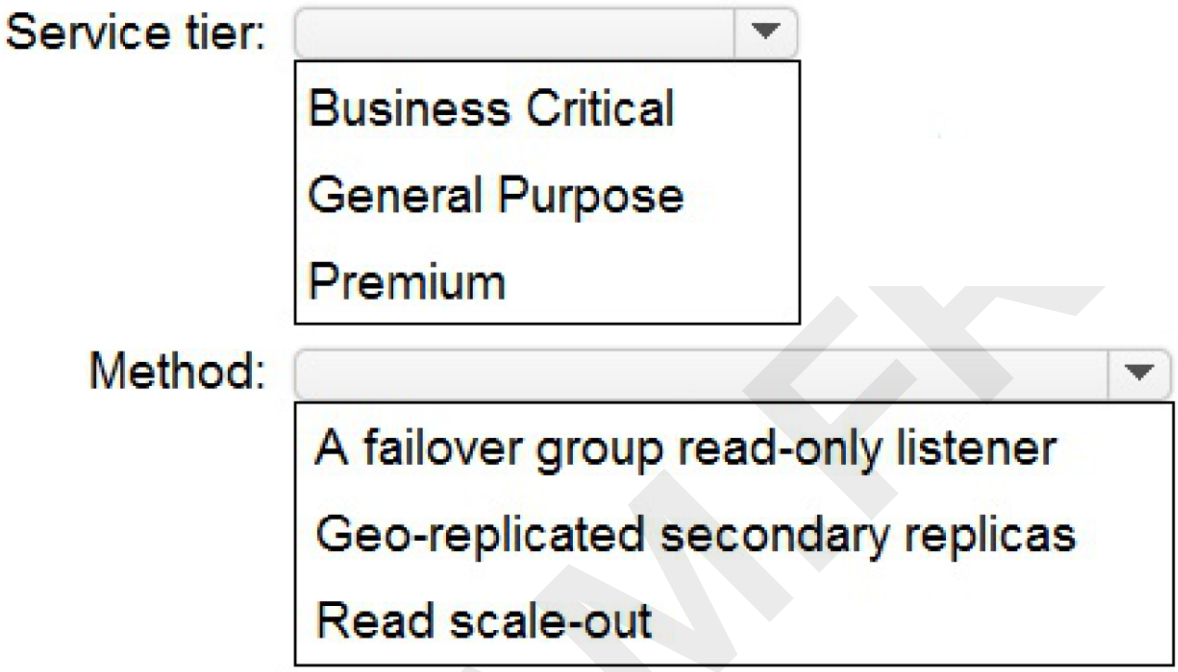

Answer:

Answer Area

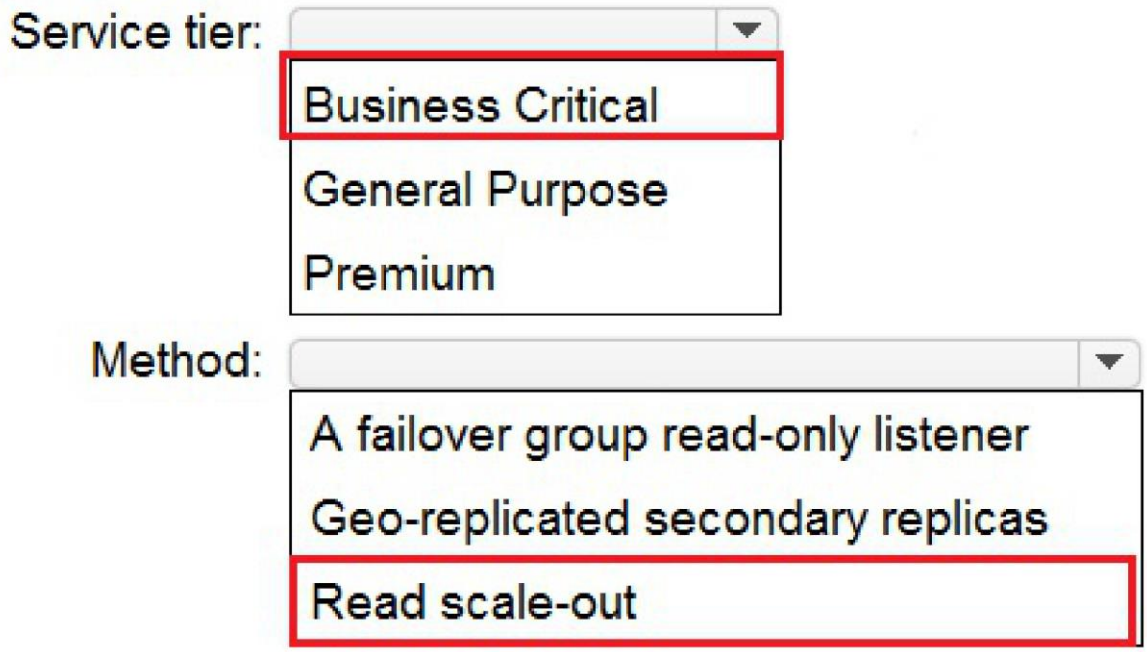

Explanation:

Service Tier: Business Critical.

Method: Read scale-out.

The read scale-out feature allows you to offload read-only workloads using the compute capacity of one of the read-only replicas, instead of running them on the read-write replica.

Read scale-out is always enabled in the Business Critical service tier of SQL Managed Instance, and for Hyperscale databases with at least one secondary replica.

https://learn.microsoft.com/en-us/azure/azure-sql/database/read-scale-out?view=azuresql

Question: 68 HOTSPOT

-

You have an on-premises Microsoft SQL Server database named DB1.

You have an Azure subscription.

You need to migrate DB1 to an Azure SQL service that meets the following requirements:
•Protects the confidentiality of sensitive data from malware and high-privileged unauthorized database administrators
•Supports pattern matching for server-side database operations
•Uses a hardware-based encryption technology
Which Azure SQL service and attestation service should you include in the solution? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

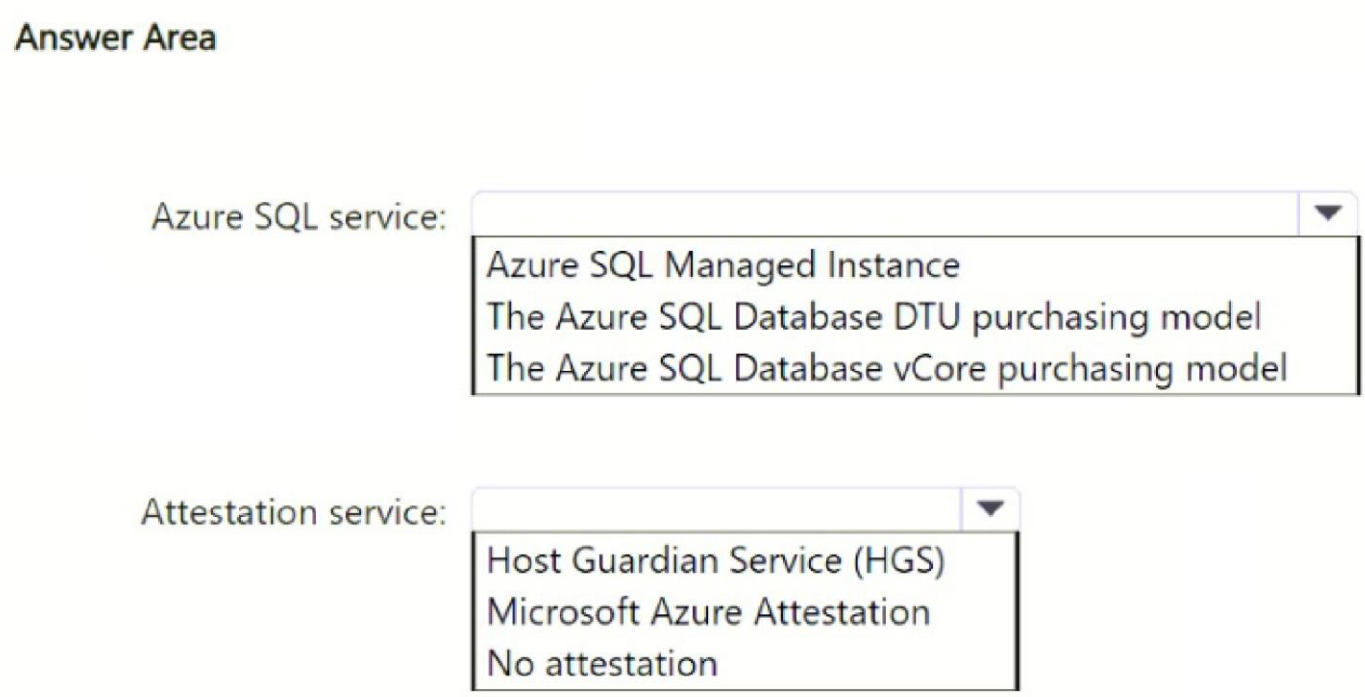

Answer:

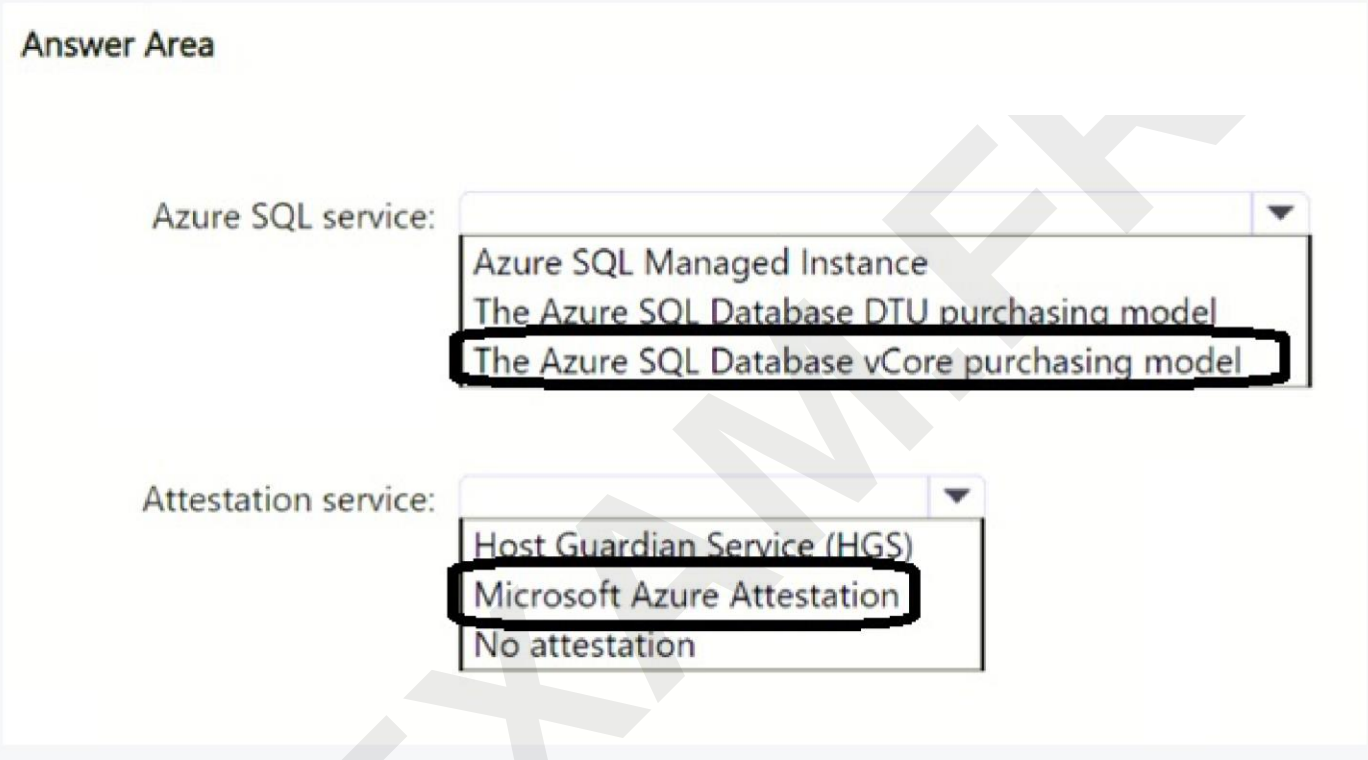

Question: 69

DRAG DROP

-

You have two on-premises Microsoft SQL Server instances named SQL1 and SQL2.

You have an Azure subscription.

You need to sync a subset of tables between the databases hosted on SQL1 and SQL2 by using SQL Data Sync.

Which five 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.

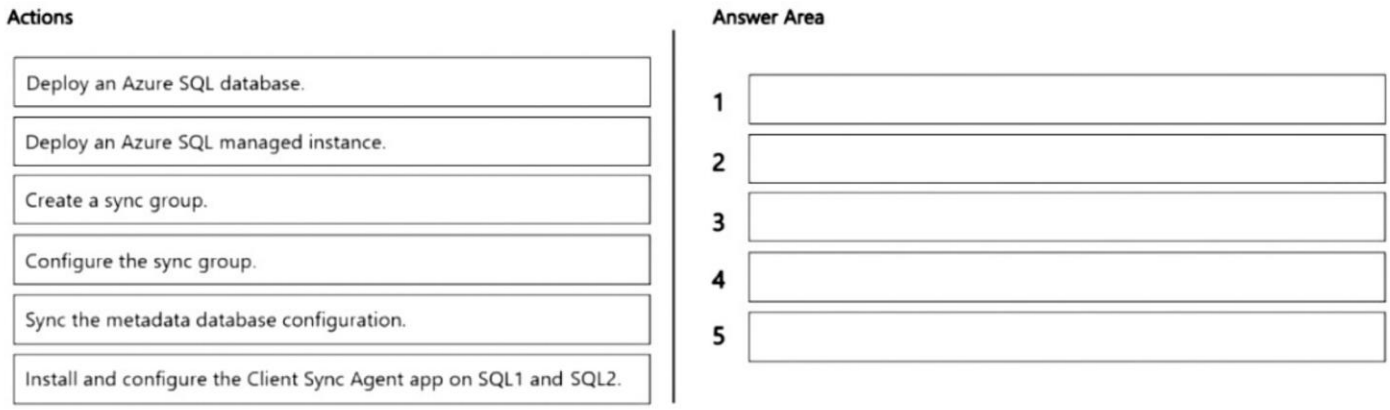

Answer:

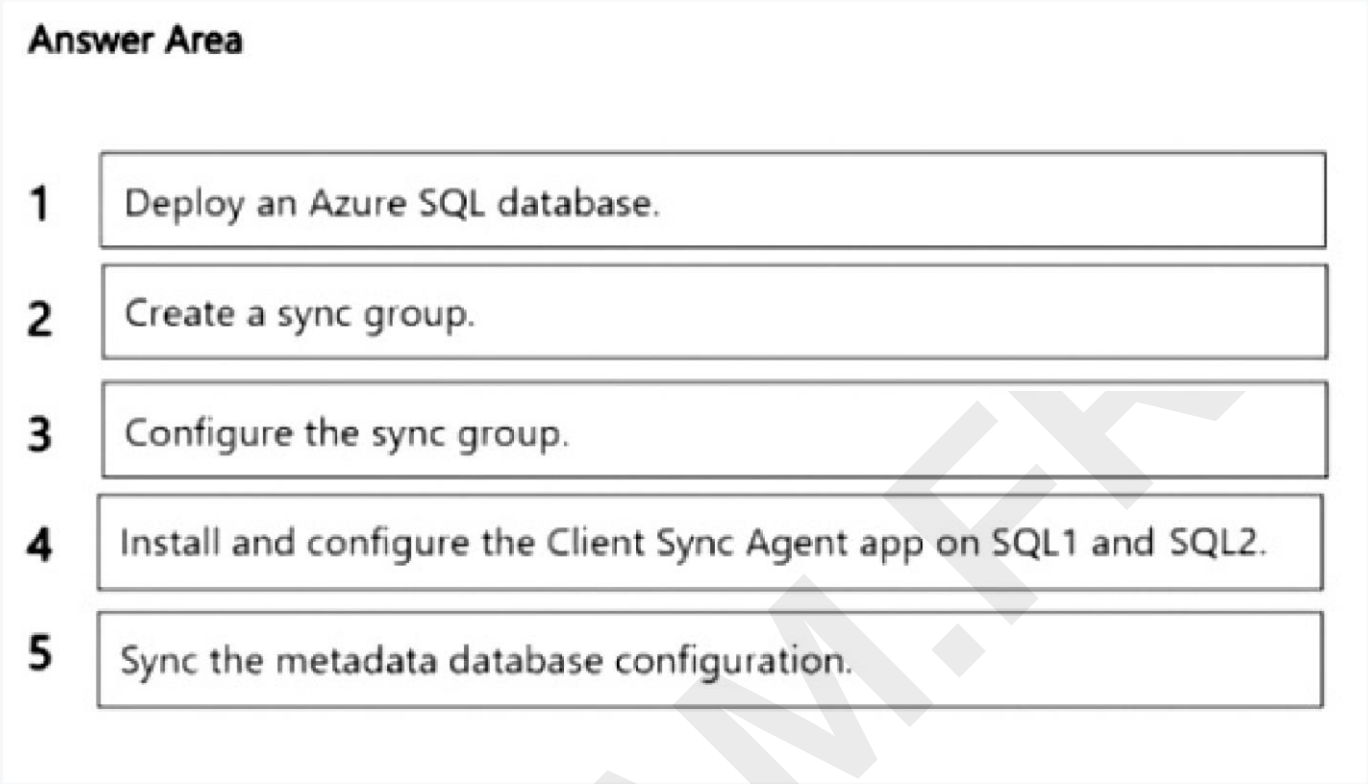

Explanation:

Deploy an Azure SQL database - This is necessary as the Azure SQL database acts as the hub database for synchronization.

Create a sync group - This will define the synchronization topology, which includes the hub and member databases.

Configure the sync group - Set up the parameters for the sync group, including the sync direction and the specific tables or columns to sync.

Install and configure the Client Sync Agent app on SQL1 and SQL2 - The Client Sync Agent enables on-premises SQL Servers to sync with Azure.

Sync the metadata database configuration - This completes the sync configuration by syncing the metadata that describes the sync relationships.

This sequence ensures that the necessary components are set up for Azure SQL Data Sync.

Question: 70 You have an on-premises Microsoft SQL Server 2022 instance that hosts a 60-TB production database named DB1. You plan to migrate DB1 to Azure.

You need to recommend a hosting solution for DB1.

Which Azure SQL Database service tier should you use to host DB1?

A.Hyperscale
B.Business Critical
C.General Purpose

Answer: A

Explanation:

Correct answer is A:Hyperscale.

Question: 71 You have an Azure subscription.

You plan to provision a single Azure SQL database.

You need to ensure that the database supports the autoscaling of compute resources. Which service tier should you choose?

A.Premium
B.General Purpose
C.Business Critical
D.Standard

Answer: B

Explanation:

Correct answer is B:General Purpose.

Question: 72 HOTSPOT -

You plan to deploy an instance of SQL Server on Linux Azure Virtual Machines. The instance will run Microsoft SQL Server 2022 and use the SQL Server IaaS Agent extension for Linux.

Which Linux operating system should you deploy, and which benefit will the SQL Server IaaS Agent extension provide? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

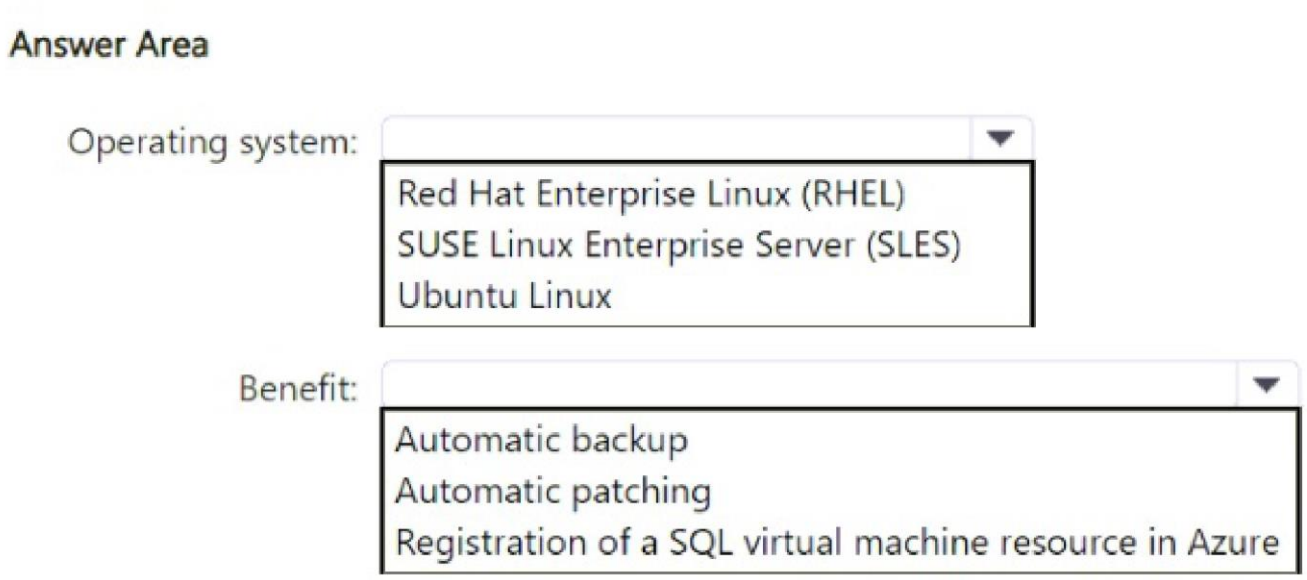

Answer:

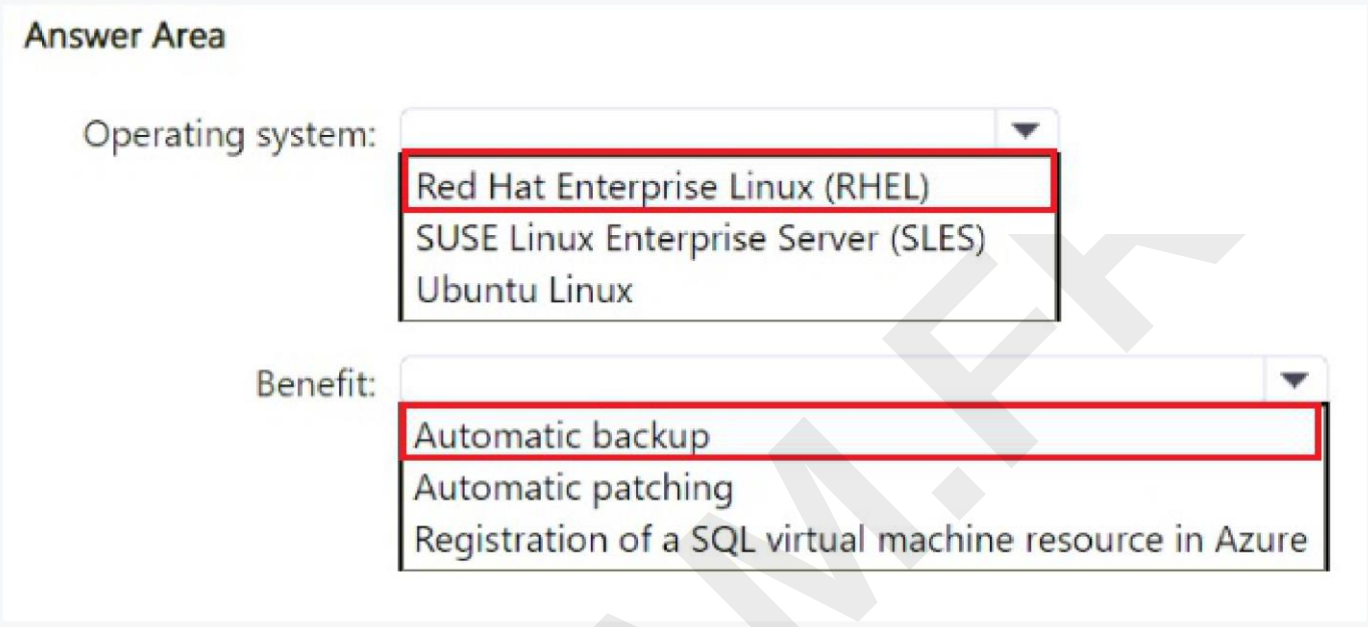

Explanation:

Operating System: Red Hat Enterprise Linux (RHEL).

Benefit: Automatic backup.

Question: 73 DRAG DROP

-

You have a burstable Azure virtual machine named VM1 that hosts an instance of Microsoft SQL Server. You need to attach an Azure ultra disk to VM1. The solution must minimize downtime on VM1.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

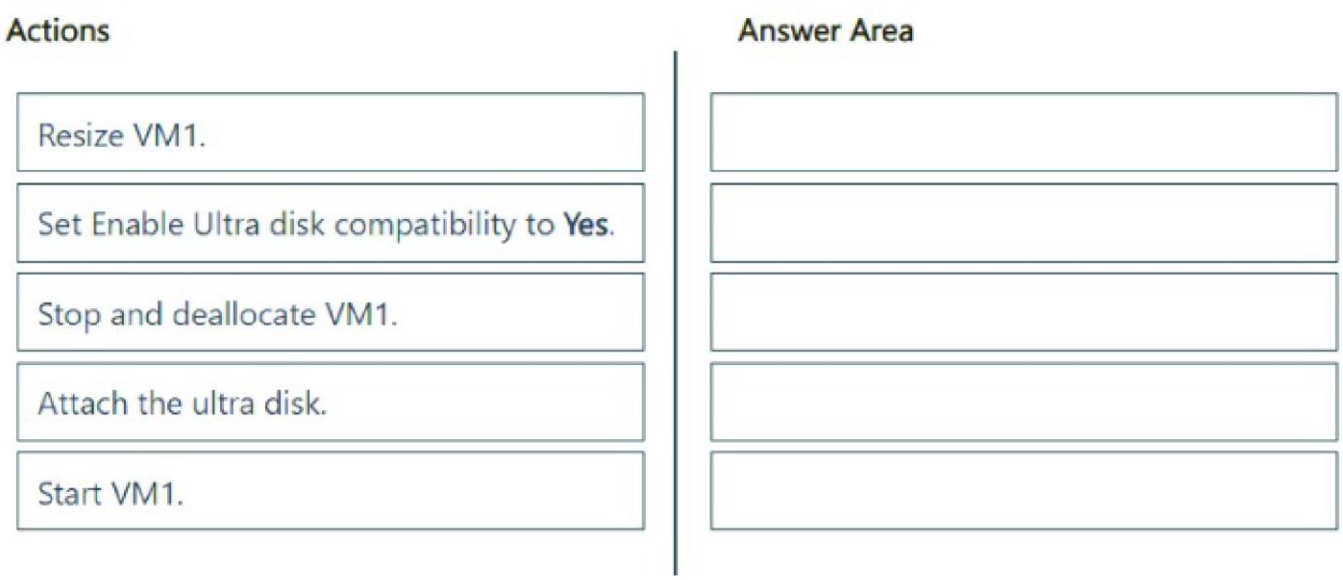

Answer:

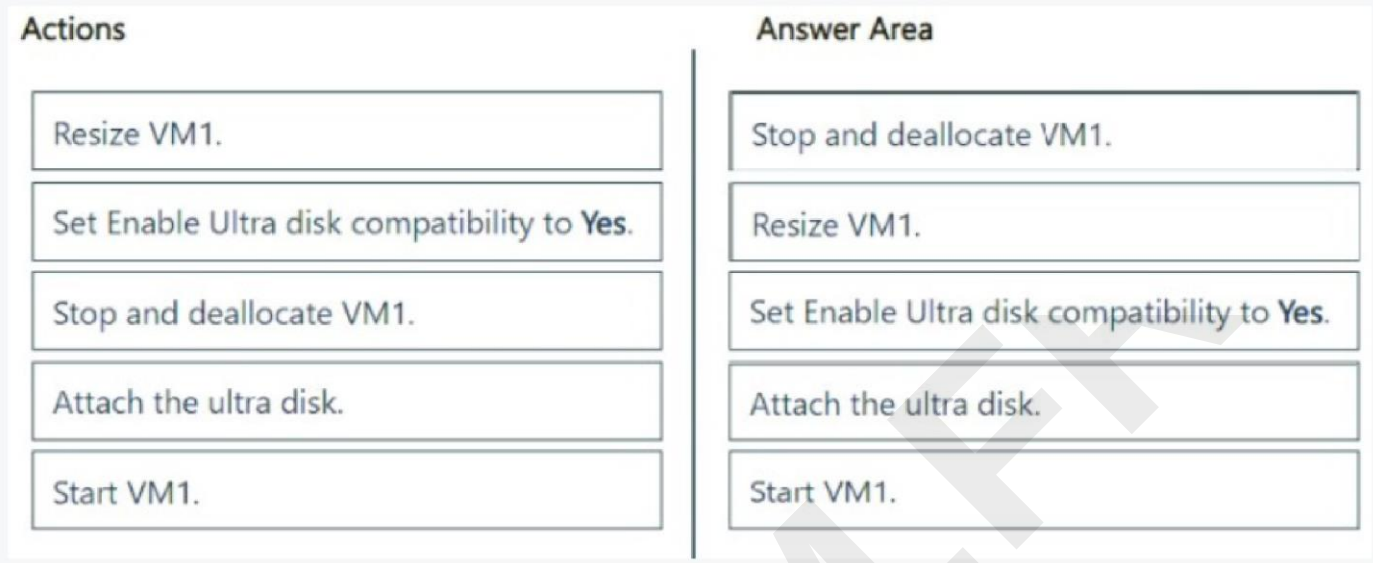

Question: 74

You have a new Azure SQL database. The database contains a column that stores confidential information. You need to track each time values from the column are returned in a query. The tracking information must be stored for 365 days from the date the query was executed.
Which three actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Turn on auditing and write audit logs to an Azure Storage account.
B. Add extended properties to the column.
C. Turn on auditing and write audit logs to an Event Hub
D. Apply sensitivity labels named Highly Confidential to the column.
E. Turn on Azure Defender for SQL

Answer: ACD

Explanation:

D: You can apply sensitivity-classification labels persistently to columns by using new metadata attributes

that have been added to the SQL Server database engine. This metadata can then be used for advanced, sensitivity-based auditing and protection scenarios.

A: An important aspect of the information-protection paradigm is the ability to monitor access to sensitive data. Azure SQL Auditing has been enhanced to include a new field in the audit log called data_sensitivity_information. This field logs the sensitivity classifications (labels) of the data that was returned by a query. Here's an example:

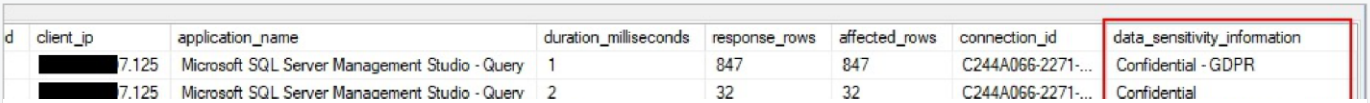

d	client_ip	application_name	duration_milliseconds	response_rows	affected_rows	connection_id	data_sensitivity_information
	[redacted]7.125	Microsoft SQL Server Management Studio - Query	1	847	847	C244A066-2271-...	Confidential - GDPR
	[redacted]7.125	Microsoft SQL Server Management Studio - Query	2	32	32	C244A066-2271-...	Confidential

C

Auditing for Azure SQL Database and Azure Synapse Analytics tracks database events and writes them to an audit log in your Azure storage account, Log Analytics workspace, or Event Hubs.

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

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/database/data-discovery-and-classification-overview
https://docs.microsoft.com/en-us/azure/azure-sql/database/azure-defender-for-sql

Question: 75

You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.
You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.
You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:
Ensure that all traffic to the public endpoint of SqlSrv1 is blocked.
Minimize the possibility of VM1 exfiltrating data stored in SqlDb1.
What should you create on VNet1?

A. a VPN gateway
B. a service endpoint
C. a private link
D. an ExpressRoute gateway

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://docs.microsoft.com/en-us/azure/private-link/private-link-overview

Question: 76

DRAG DROP -

You have a new Azure SQL database named DB1 on an Azure SQL server named AzSQL1.
The only user who was created is the server administrator.
You need to create a contained database user in DB1 who will use Azure Active Directory (Azure AD) for authentication.
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:

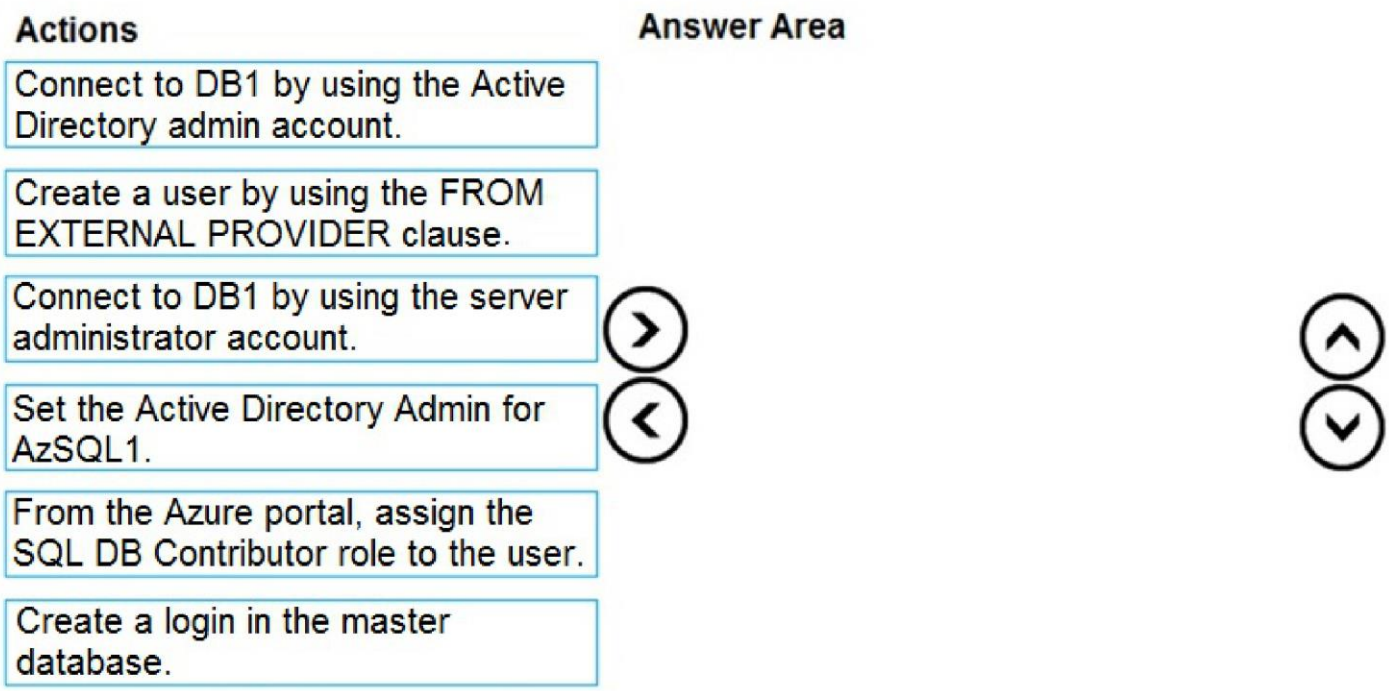

Answer:

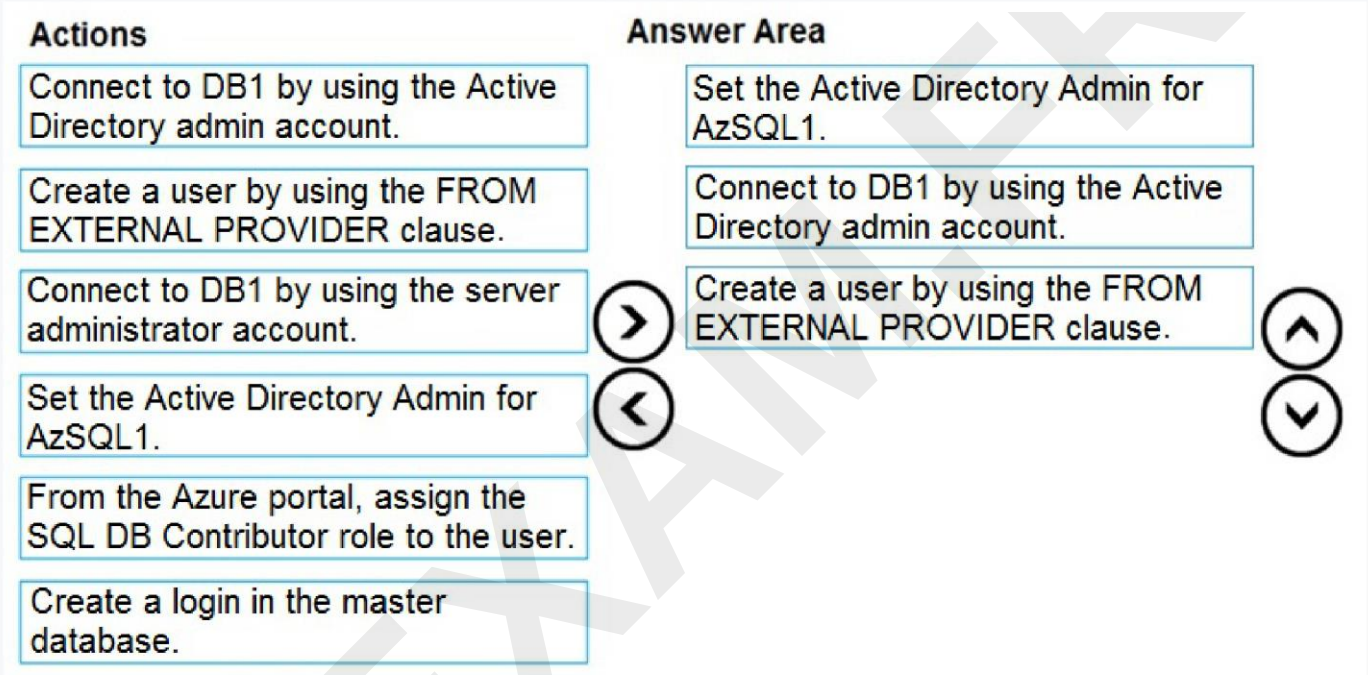

Explanation:

Step 1: Set up the Active Directory Admin for AzSQL1.

Step 2: Connect to DB1 by using the Active Directory admin account.

Step 3: Create a user by using the FROM EXTERNAL PROVIDER clause.

FROM EXTERNAL PROVIDER is available for creating server-level Azure AD logins in SQL Database managed instance. Azure AD logins allow database-level

Azure AD principals to be mapped to server-level Azure AD logins. To create an Azure AD user from an Azure AD login use the following syntax:

CREATE USER [AAD_principal] FROM LOGIN [Azure AD login]

Reference:

https://docs.microsoft.com/en-us/sql/t-sql/statements/create-user-transact-sql

Question: 77

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

Customer_ID	Customer_Name	Customer_Phone
11001	Contoso, Ltd.	555-555-0173
11002	Litware, Inc.	555-505-3124
11003	ADatum Corporation	555-689-4312

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

- The first six numerals of each customer's phone number must be masked.
- The last four digits of each customer's phone number must be visible.
- Hyphens must be preserved and displayed.

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

Answer Area

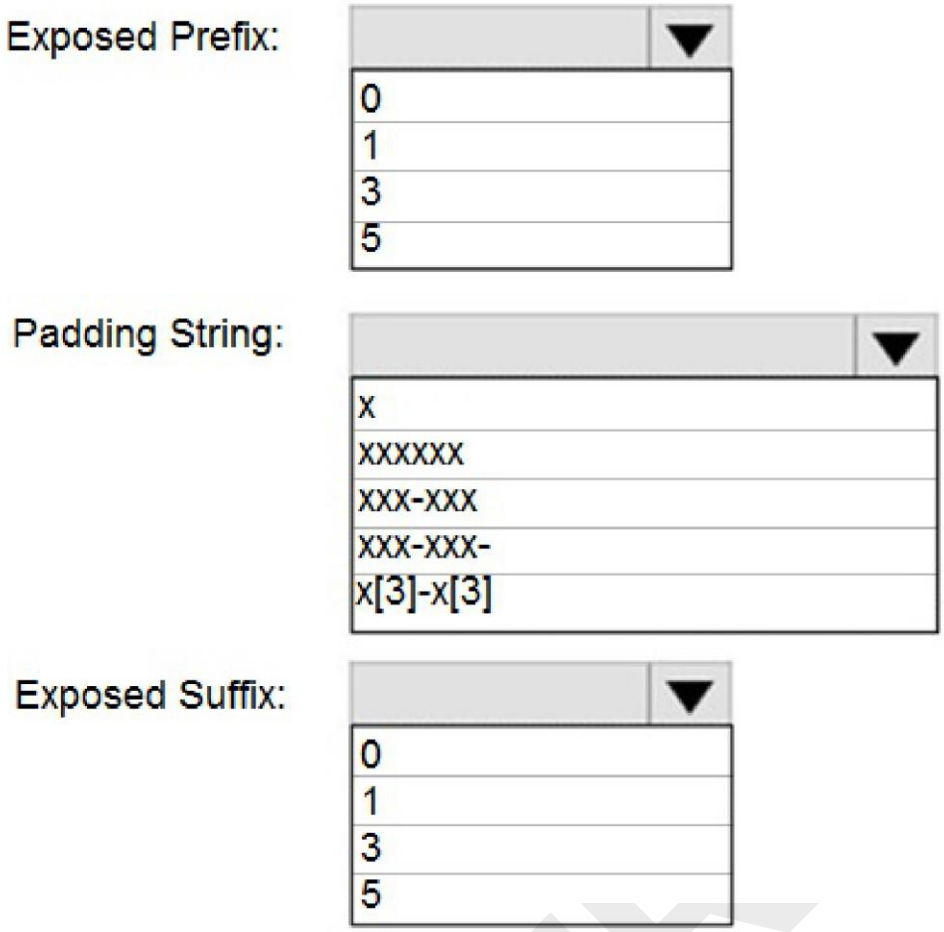

Answer:

Answer Area

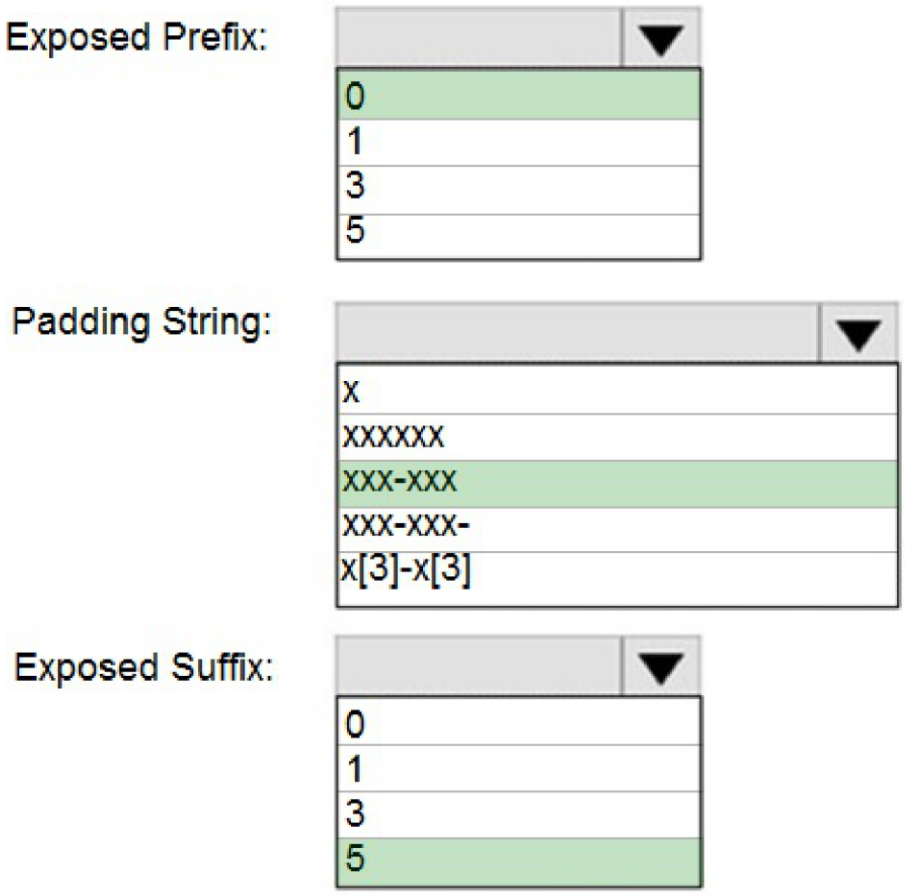

Explanation:
Box 1: 0 -
Custom String : Masking method that exposes the first and last letters and adds a custom padding string in the middle. prefix,[padding],suffix

Box 2: xxx-xxx -

Box 3: 5 -

Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking

Question: 78

DRAG DROP -
You have an Azure SQL database that contains a table named Employees. Employees contains a column named Salary.
You need to encrypt the Salary column. The solution must prevent database administrators from reading the data in the Salary column and must provide the most secure encryption.
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:

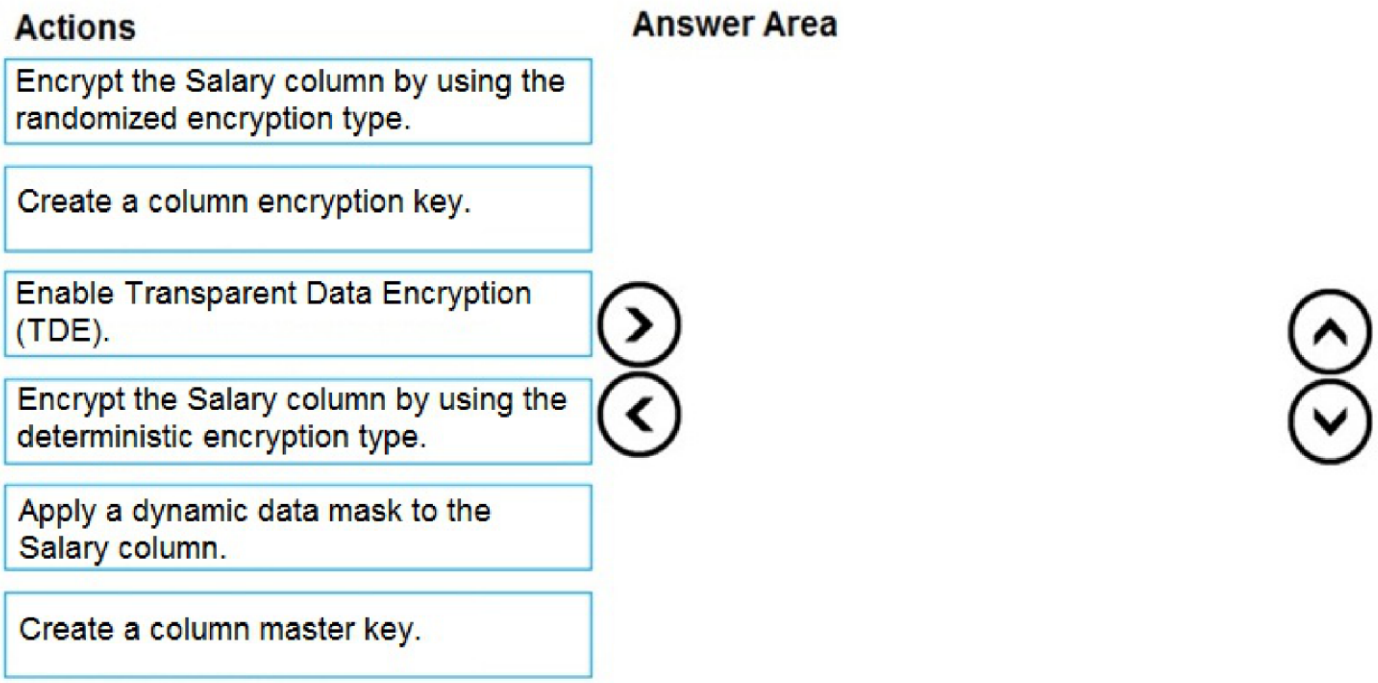

Answer:

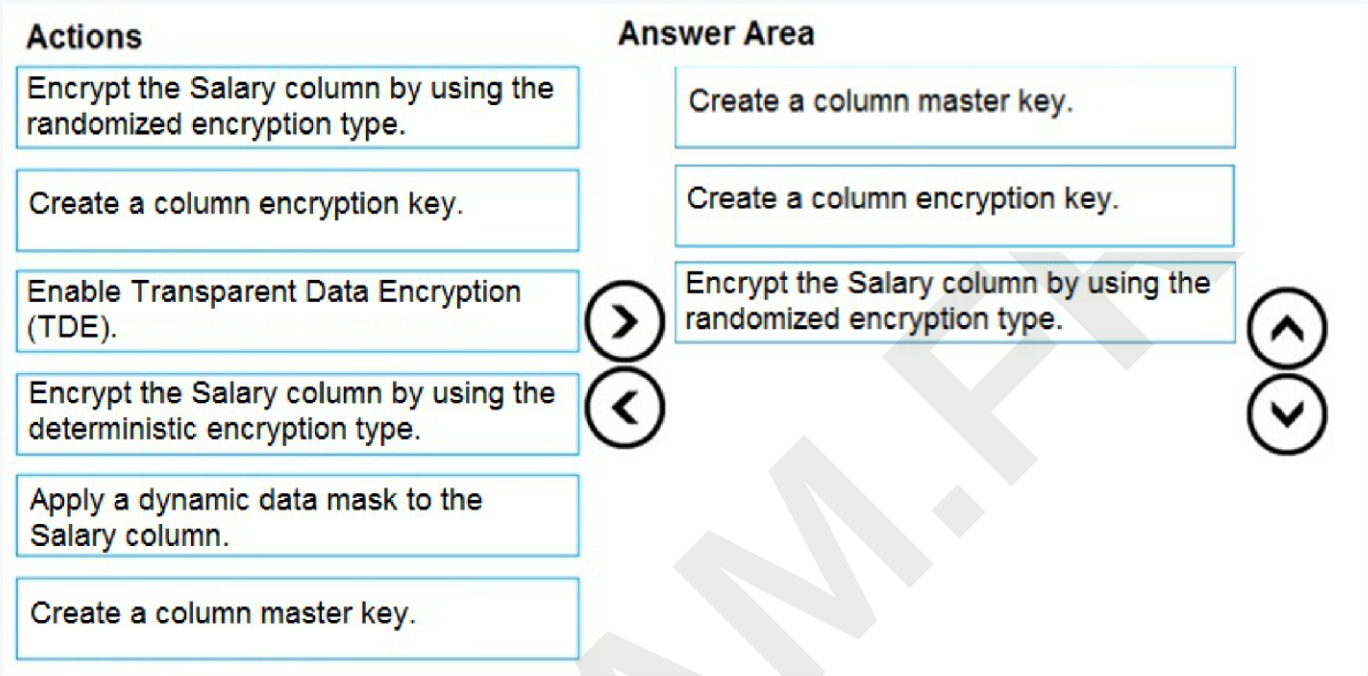

Explanation:

Step 1: Create a column master key

Create a column master key metadata entry before you create a column encryption key metadata entry in the database and before any column in the database can be encrypted using Always Encrypted.

Step 2: Create a column encryption key.

Step 3: Encrypt the Salary column by using the randomized encryption type.

Randomized encryption uses a method that encrypts data in a less predictable manner. Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns. Note: A column encryption key metadata object contains one or two encrypted values of a column encryption key that is used to encrypt data in a column. Each value is encrypted using a column master key.

Incorrect Answers:

Deterministic encryption.

Deterministic encryption always generates the same encrypted value for any given plain text value. Using deterministic encryption allows point lookups, equality joins, grouping and indexing on encrypted columns.

However, it may also allow unauthorized users to guess information about encrypted values by examining

patterns in the encrypted column, especially if there's a small set of possible encrypted values, such as True/False, or North/South/East/West region.

Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine

Question: 79

HOTSPOT -
You have an Azure SQL database named DB1 that contains two tables named Table1 and Table2. Both tables contain a column named a Column1. Column1 is used for joins by an application named App1.
You need to protect the contents of Column1 at rest, in transit, and in use.
How should you protect the contents of Column1? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.
Hot Area:

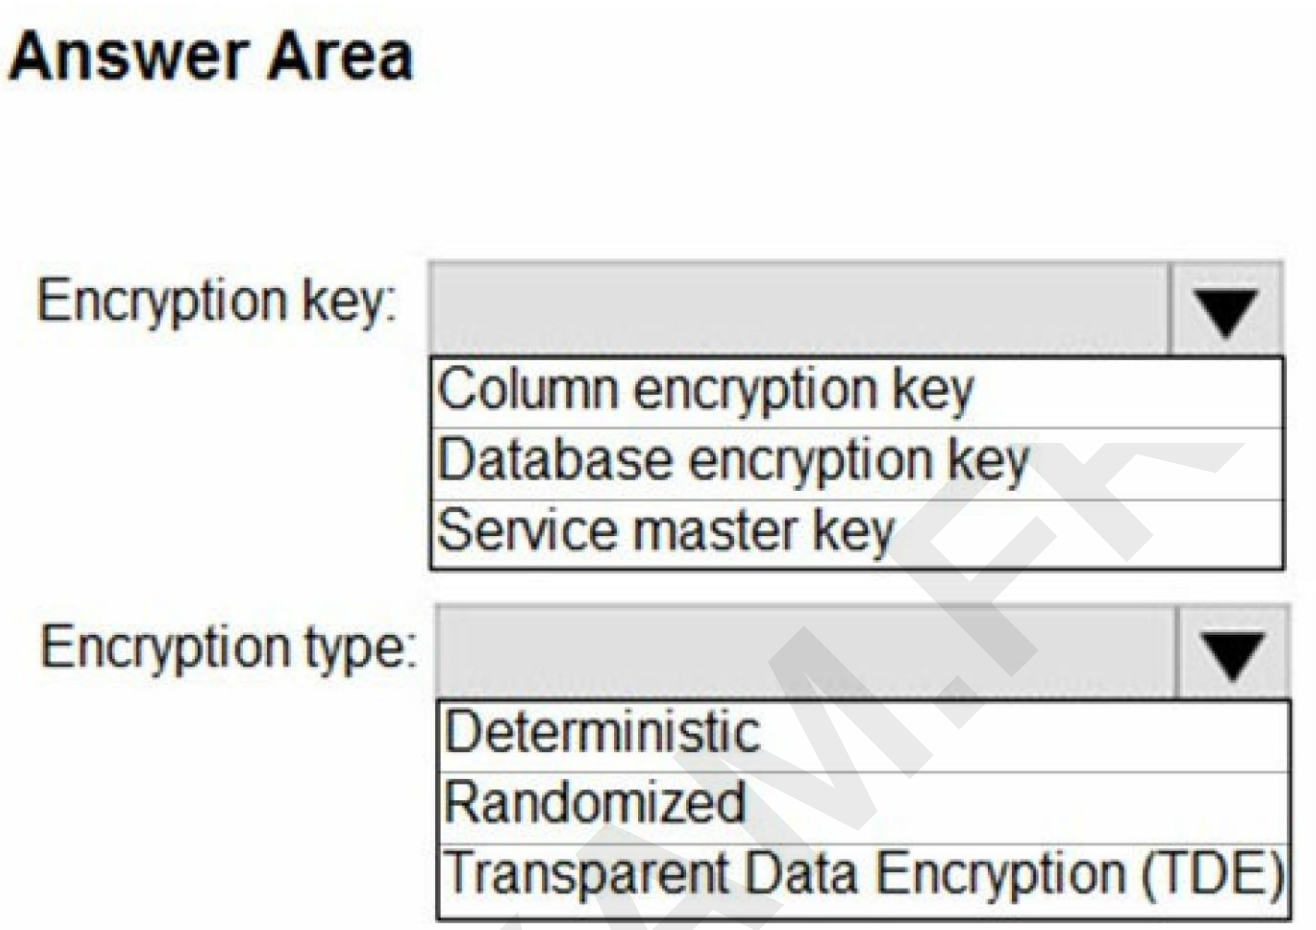

Answer:

Answer Area

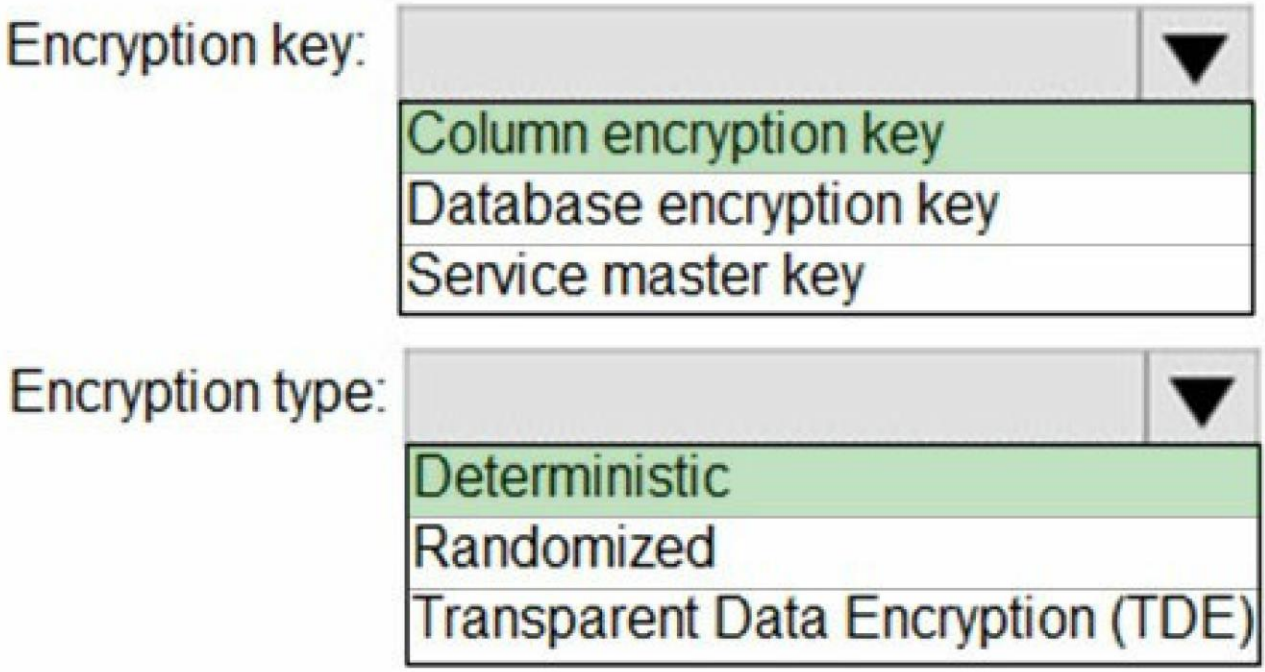

Explanation:
Box 1: Column encryption Key -
Always Encrypted uses two types of keys: column encryption keys and column master keys. A column encryption key is used to encrypt data in an encrypted column. A column master key is a key-protecting key that encrypts one or more column encryption keys.

Incorrect Answers:
TDE encrypts the storage of an entire database by using a symmetric key called the Database Encryption Key (DEK).

Box 2: Deterministic -
Always Encrypted is a feature designed to protect sensitive data, such as credit card numbers or national identification numbers (for example, U.S. social security numbers), stored in Azure SQL Database or SQL Server databases. Always Encrypted allows clients to encrypt sensitive data inside client applications and never reveal the encryption keys to the Database Engine (SQL Database or SQL Server).

Always Encrypted supports two types of encryption: randomized encryption and deterministic encryption.

Deterministic encryption always generates the same encrypted value for any given plain text value. Using deterministic encryption allows point lookups, equality joins, grouping and indexing on encrypted columns.

Incorrect Answers:

☞ Randomized encryption uses a method that encrypts data in a less predictable manner. Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns.

☞ Transparent data encryption (TDE) helps protect Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics against the threat of malicious offline activity by encrypting data at rest. It performs real-time encryption and decryption of the database, associated backups, and transaction log files at rest without requiring changes to the application.

Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-e ngine

Question: 80

You have 40 Azure SQL databases, each for a different customer. All the databases reside on the same Azure SQL Database server.
You need to ensure that each customer can only connect to and access their respective database.
Which two actions should you perform? Each correct answer presents part of the solution. NOTE:
Each correct selection is worth one point.

A. Implement row-level security (RLS).
B. Create users in each database.
C. Configure the database firewall.
D. Configure the server firewall.
E. Create logins in the master database.
F. Implement Always Encrypted.

Answer: BC

Explanation:

On Azure SQL databases, you do not have to create a Login for the user unless he needs to access several databases within the server. The requirement is each user will only be allowed to access his own database. So with the principle of least privilege, we just need to a contained user. The database firewall is additional security to limit access on each database.