

➤ **Vendor:** Microsoft

➤ **Exam Code:** DP-200

➤ **Exam Name:** Implementing an Azure Data Solution

➤ **New Updated Questions from** [Braindump2go](#) (Updated in [May/2020](#))

[Visit Braindump2go and Download Full Version DP-200 Exam Dumps](#)

**QUESTION 164**

**SIMULATION**

Use the following login credentials as needed:

Azure Username: xxxxx

Azure Password: xxxxx

The following information is for technical support purposes only:

Lab Instance: 10543936



Your company's security policy states that administrators must be able to review a list of the failed logins to an Azure SQL database named db1 during the previous 30 days.

You need to modify your Azure environment to meet the security policy requirements.

**To complete this task, sign in to the Azure portal.**

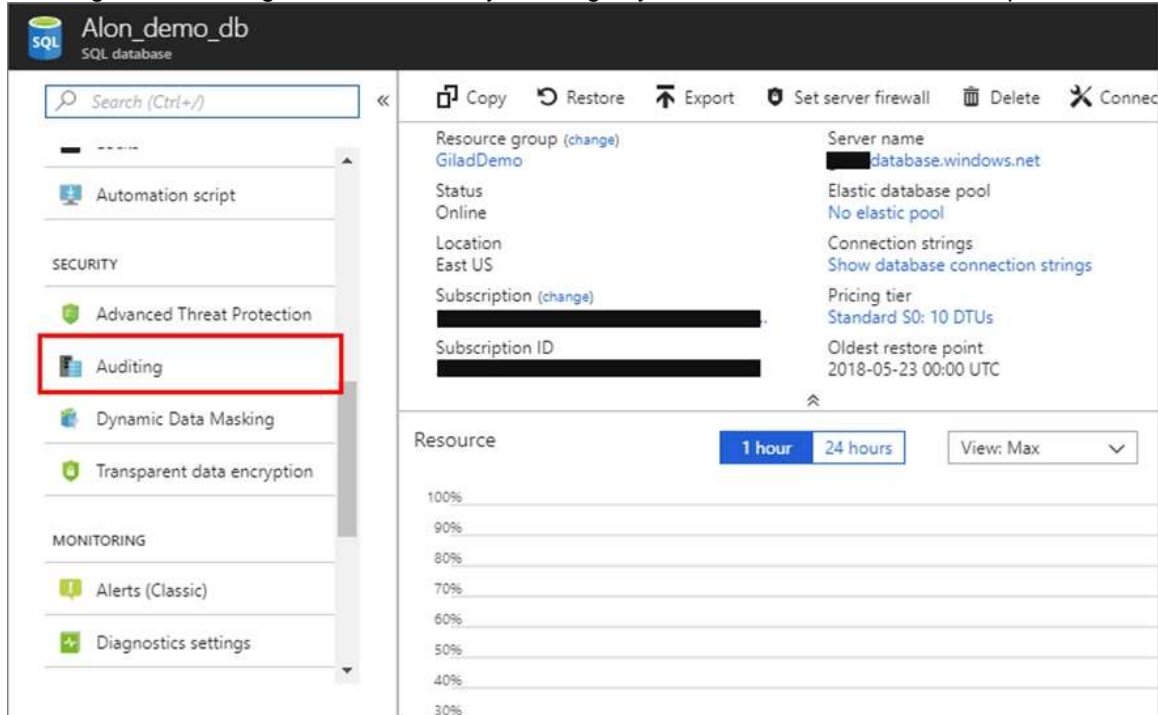
**Answer:**

Set up auditing for your database

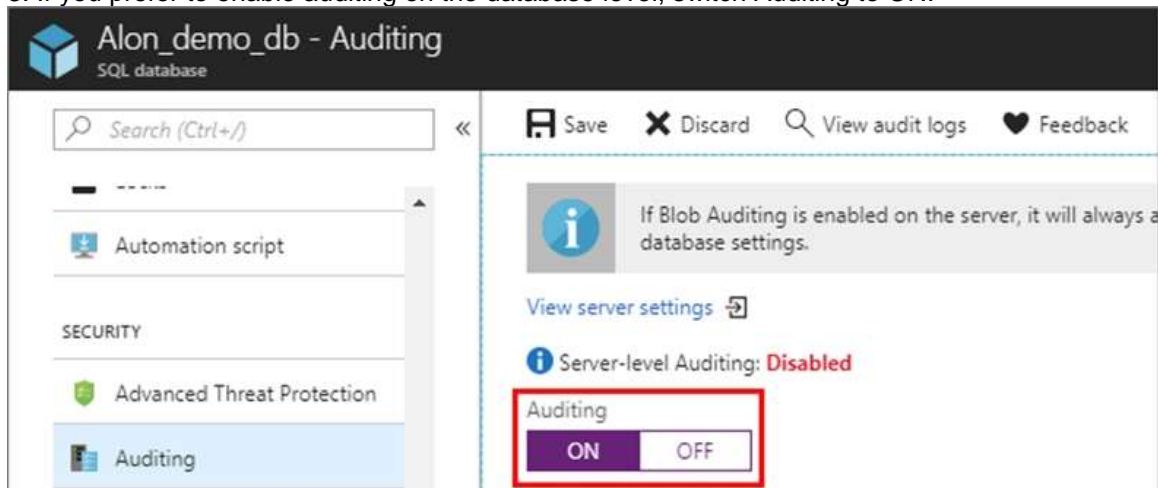
The following section describes the configuration of auditing using the Azure portal.

1. Go to the Azure portal.

2. Navigate to Auditing under the Security heading in your SQL database db1/server pane



3. If you prefer to enable auditing on the database level, switch Auditing to ON.



Reference:

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

## QUESTION 165

### SIMULATION

Use the following login credentials as needed:

Azure Username: xxxxx

Azure Password: xxxxx

The following information is for technical support purposes only:

Lab Instance: 10543936



You need to ensure that all REST API calls to an Azure Storage account named storage10543936 use HTTPS only.

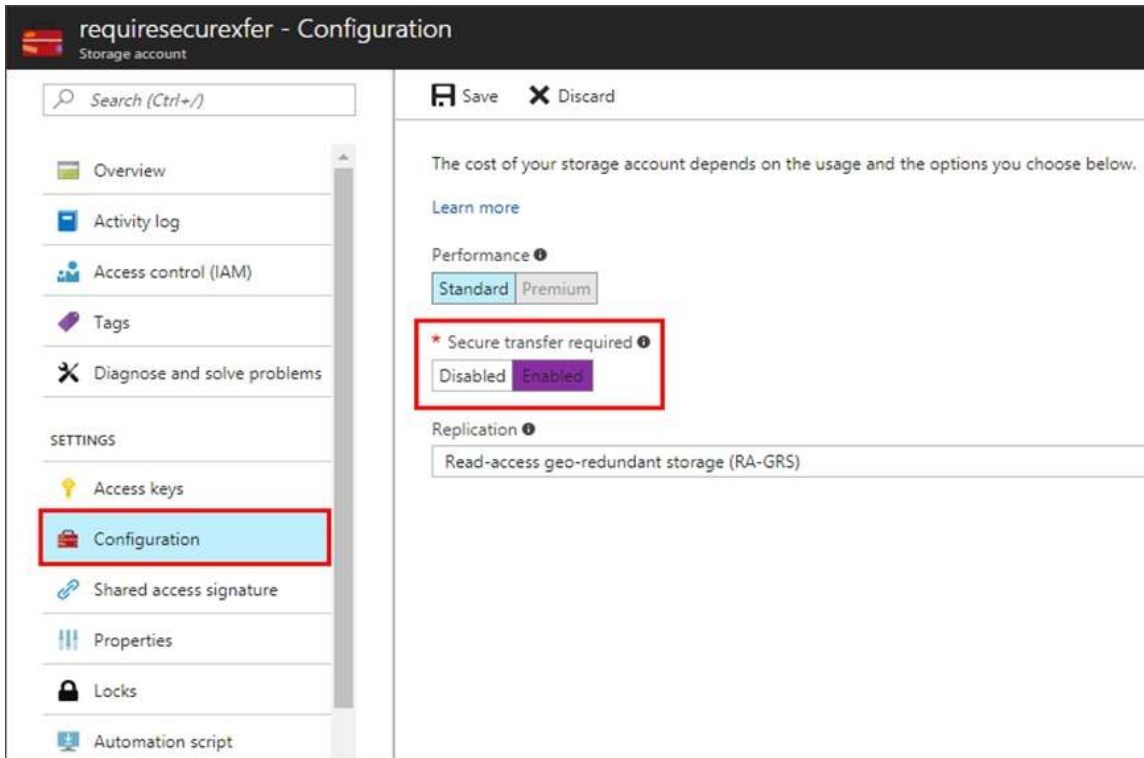
**To complete this task, sign in to the Azure portal.**

**Answer:**

You can configure your storage account to accept requests from secure connections only by setting the Secure transfer required property for the storage account.

Require secure transfer for an existing storage account

1. Select the existing storage account storage10543936 in the Azure portal.
2. In the storage account menu pane, under SETTINGS, select Configuration.
3. Under Secure transfer required, select Enabled.



Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-require-secure-transfer>

#### QUESTION 166

##### SIMULATION

Use the following login credentials as needed:

Azure Username: xxxxx

Azure Password: xxxxx

The following information is for technical support purposes only:

Lab Instance: 10543936



You need to double the available processing resources available to an Azure SQL data warehouse named datawarehouse.

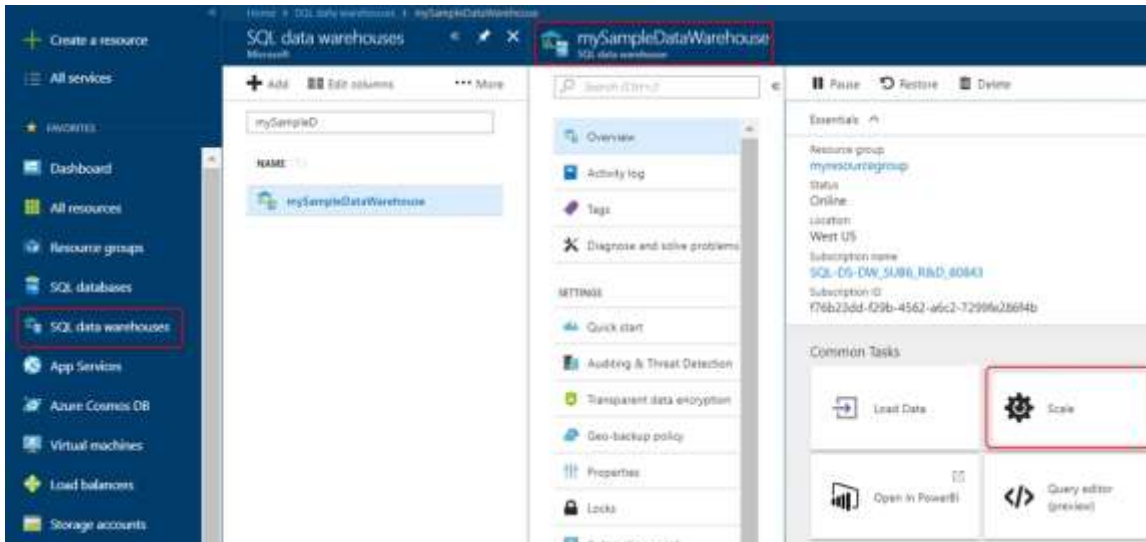
**To complete this task, sign in to the Azure portal.**

**NOTE: This task might take several minutes to complete. You can perform other tasks while the task completes or end this section of the exam.**

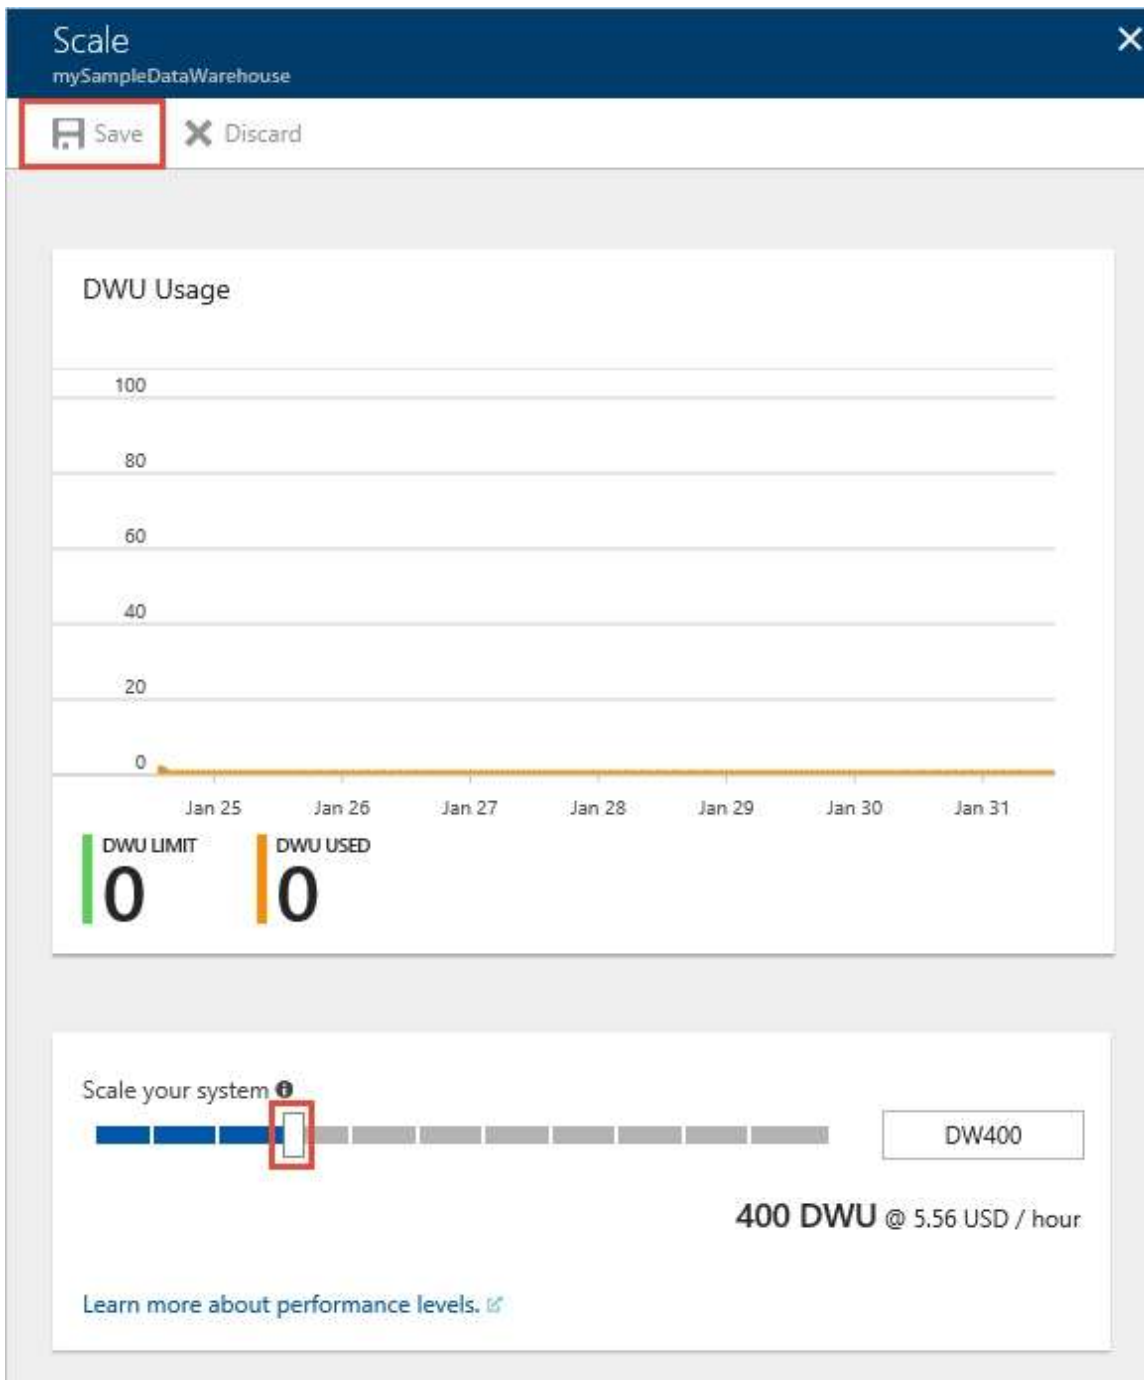
**Answer:**

SQL Data Warehouse compute resources can be scaled by increasing or decreasing data warehouse units.

1. Click SQL data warehouses in the left page of the Azure portal.
2. Select datawarehouse from the SQL data warehouses page. The data warehouse opens.
3. Click Scale.





4. In the Scale panel, move the slider left or right to change the DWU setting. Double the DWU setting.



6. Click Save. A confirmation message appears. Click yes to confirm or no to cancel.



**Scale** mySampleDataWarehouse ×


 **Save**  **Discard**

### Save performance level

Scaling the performance level of your SQL Data Warehouse will cause all inflight queries to be canceled which may result in minutes to hours of downtime until rollback is complete. To learn more about monitoring your queries, please visit our online documentation for monitoring. We do not recommend using scale to abort long running queries.

**Yes**

No




DWU LIMIT

0

DWU USED

0

Scale your system ⓘ



DW500

**500 DWU @ 6.96 USD / hour**

[Learn more about performance levels.](#)

Reference:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/quickstart-scale-compute-portal>**QUESTION 167**

Hotspot Question

You have the following Azure Stream Analytics query.



WITH

```
step1 AS (SELECT *  
  FROM input1  
  PARTITION BY StateID  
  INTO 10),  
step2 AS (SELECT *  
  FROM input2  
  PARTITION BY StateID  
  INTO 10)
```

```
SELECT *  
  INTO output  
  FROM step1  
  PARTITION BY StateID  
  UNION step2  
  BY StateID
```

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

### Answer Area

Statements	Yes	No
The query joins two streams of partitioned data.	<input type="radio"/>	<input type="radio"/>
The stream scheme key and count must match the output scheme.	<input type="radio"/>	<input type="radio"/>
Providing 60 streaming units will optimize the performance of the query.	<input type="radio"/>	<input type="radio"/>

Answer:

### Answer Area

Statements	Yes	No
The query joins two streams of partitioned data.	<input checked="" type="radio"/>	<input type="radio"/>
The stream scheme key and count must match the output scheme.	<input checked="" type="radio"/>	<input type="radio"/>
Providing 60 streaming units will optimize the performance of the query.	<input checked="" type="radio"/>	<input type="radio"/>

#### Explanation:

Box 1: Yes

You can now use a new extension of Azure Stream Analytics SQL to specify the number of partitions of a stream when reshuffling the data.

The outcome is a stream that has the same partition scheme. Please see below for an example:

WITH step1 AS (SELECT \* FROM [input1] PARTITION BY DeviceID INTO 10), step2 AS (SELECT \* FROM [input2] PARTITION BY DeviceID INTO 10)

SELECT \* INTO [output] FROM step1 PARTITION BY DeviceID UNION step2 PARTITION BY DeviceID

Note: The new extension of Azure Stream Analytics SQL includes a keyword INTO that allows you to specify the number of partitions for a stream when performing reshuffling using a PARTITION BY statement.

Box 2: Yes

When joining two streams of data explicitly repartitioned, these streams must have the same partition key and partition count.

Box 3: Yes

10 partitions x six SUs = 60 SUs is fine.

Note: Remember, Streaming Unit (SU) count, which is the unit of scale for Azure Stream Analytics, must be adjusted so the number of physical resources available to the job can fit the partitioned flow. In general, six SUs is a good number to assign to each partition. In case there are insufficient resources assigned to the job, the system will only apply the repartition if it benefits the job.

Reference:

<https://azure.microsoft.com/en-in/blog/maximize-throughput-with-repartitioning-in-azure-stream-analytics/>

#### QUESTION 168

A company has an Azure SQL Database defined as part of their Azure subscription. The Automatic tuning settings are configured as shown below:

Option	Server level	Database level
Force Plan	Inherited	Inherited
Create Index	Inherited	Inherited
Drop Index	Inherited	Inherited

Would the setting of "Create Index" be ON for the database?

- A. Yes
- B. No

**Answer: A**

#### Explanation:

If you implement the below auditing settings for the server

Inherit from: ⓘ

**Azure defaults** Don't inherit

ⓘ The database is inheriting automatic tuning configuration from Azure defaults.

Configure the automatic tuning options ⓘ

OPTION	DESIRED STATE	CURRENT STATE
 FORCE PLAN	ON OFF <b>INHERIT</b>	<b>ON</b> Inherited from Azure defaults
 CREATE INDEX	ON OFF <b>INHERIT</b>	<b>ON</b> Inherited from Azure defaults
 DROP INDEX	ON OFF <b>INHERIT</b>	<b>OFF</b> Inherited from Azure defaults

And also place the following settings for the database

Inherit from: ⓘ

**Server** Azure defaults Don't inherit

ⓘ The database is inheriting automatic tuning configuration from the server. You can set the configuration to be inherited by going [tuning settings](#)

Configure the automatic tuning options ⓘ

OPTION	DESIRED STATE	CURRENT STATE
 FORCE PLAN	ON OFF <b>INHERIT</b>	<b>ON</b> Inherited from server
 CREATE INDEX	ON OFF <b>INHERIT</b>	<b>ON</b> Inherited from server
 DROP INDEX	ON OFF <b>INHERIT</b>	<b>OFF</b> Inherited from server

You can see that the “CREATE INDEX” setting is ON.

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-automatic-tuning-enable>

### QUESTION 169

You need to create a new Azure Databricks cluster. This cluster would connect to Azure Data Lake Storage Gen2 by using Azure Active Directory (Azure AD) integration.

Which of the following advanced option would you enable?

- A. Blob access control
- B. Table access control
- C. Credential Passthrough
- D. Single Sign-On

**Answer: C**

**Explanation:**

The documentation for Azure Data bricks mentions that you should set Credential Passthrough in the Advanced option.

## Enable Azure Data Lake Storage credential passthrough for a high-concurrency cluster

High concurrency clusters can be shared by multiple users. They support only Python, SQL, and R.

1. When you create a cluster, set the Cluster Mode to **High Concurrency**.
2. Choose a Databricks Runtime version according to the Azure Data Lake Storage type:
  - Azure Data Lake Storage Gen1: Databricks Runtime 5.1 or above.
  - Azure Data Lake Storage Gen2: Databricks Runtime 5.3 or above.
3. Under Advanced Options, select **Enable credential passthrough and only allow Python and SQL commands**.



<https://docs.azuredatabricks.net/data/data-sources/azure/adls-passthrough.html>

### QUESTION 170

You have JSON files stored in an Azure Data Lake Storage Gen2 account. The JSON file contains the FirstName and LastName of customers. You need to use Azure Data bricks to copy the data in the JSON files to an Azure data warehouse. A new column must be created which concatenates the FirstName and LastName values. You have the following components in place in Azure:

- A destination table in the SQL Data Warehouse
- An Azure Blob storage container
- A service principal

Which of the following are actions you would perform to transfer the data onto the Azure SQL Data warehouse table? (Choose 5)

- A. Write the results onto Azure Data Lake Storage
- B. Drop the data frame
- C. Perform transformations on the data frame
- D. Mount the Data Lake Storage onto DBFS
- E. Perform transformations on the file
- F. Read the file into a data frame
- G. Specify a temporary folder to stage the data
- H. Write the results to a table in SQL Data Warehouse

**Answer:** CDFGH

#### Explanation:

Option A is incorrect since you don't need to write the results onto Azure Data Lake storage. The results are written to the Azure SQL Data warehouse table.

Option B is incorrect since we don't need to drop the data frames.

Option B is incorrect since transformations need to be carried out on the data frames.

<https://docs.microsoft.com/en-us/azure/azure-databricks/databricks-extract-load-sql-data-warehouse#load-data-into-azure-sql-data-warehouse>

### QUESTION 171

You have an Azure SQL Database named lead2pass-db. The database contains a table named lead2pass-customer. The table has a column named customerID that is of the type varchar(22). You have to implement masking for the customerID which would meet the following requirements:

- The first two prefix characters must be exposed
- The last four prefix characters must be exposed
- All other characters must be masked

You decide to implement data masking and use a random number function mask.

Would this fulfil the requirement?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

This masking function is used to mask specific boundaries.  
The Microsoft documentation mentions the following on the function mask.

**Random number**      **Masking method, which generates a random number** according to the selected boundaries and actual data types. If the designated boundaries are equal, then the masking function is a constant number.

Masking Field Format

Random number ▼

From      To

0      0

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dynamic-data-masking-get-started>

**QUESTION 172**

A team currently managed Azure HDInsight cluster. The team spends quite a lot of time on creating and destroying clusters. They want to implement a solution that can be used to deploy Azure HDInsight clusters with minimal effort. Which of the following can they implement for this requirement?

- A. Azure Databricks
- B. Azure Load Balancer
- C. Azure Resource Manager templates
- D. Azure SQL data warehouse

**Answer: C****Explanation:**

You can use Azure Resource Manager templates to easily deploy resources in Azure. If you have to repeatable deploy the same infrastructure on Azure, you can make use of Azure Resource Manager templates.

Option A is incorrect since this is a separate Spark-based analytics platform.

Option B is incorrect since this is used to distribute traffic to backend virtual machines.

Option D is incorrect since this is data warehousing solution available on the Azure platform.

<https://docs.microsoft.com/en-us/azure/templates/microsoft.hdinsight/2018-06-01-preview/clusters>

**QUESTION 173**

A company needs to configure synchronization of data between their on-premise Microsoft SQL Server database and Azure SQL database. The synchronization process must include the following:

- Be able to perform an initial data synchronization to the Azure SQL Database with minimal downtime.
- Be able to perform bi-directional synchronization after the initial synchronization is complete

Which of the following would you consider as the synchronization solution?

- A. Data Migration Assistant
- B. Backup and restore
- C. SQL Server Agent Job
- D. Azure SQL Data Sync

**Answer: D****Explanation:**

Azure SQL Data Sync can be used to synchronize data between the on-premise SQL Server and the Azure SQL database.



## When to use Data Sync

Data Sync is useful in cases where data needs to be kept updated across several Azure SQL databases or SQL Server databases. Here are the main use cases for Data Sync:

- **Hybrid Data Synchronization:** With Data Sync, you can keep data synchronized between your on-premises databases and Azure SQL databases to enable hybrid applications. This capability may appeal to customers who are considering moving to the cloud and would like to put some of their application in Azure.
- **Distributed Applications:** In many cases, it's beneficial to separate different workloads across different databases. For example, if you have a large production database, but you also need to run a reporting or analytics workload on this data, it's helpful to have a second database for this additional workload. This approach minimizes the performance impact on your production workload. You can use Data Sync to keep these two databases synchronized.
- **Globally Distributed Applications:** Many businesses span several regions and even several countries/regions. To minimize network latency, it's best to have your data in a region close to you. With Data Sync, you can easily keep databases in regions around the world synchronized.

Option A is incorrect since this is just used to assess databases for the migration process.

Option B is incorrect since this would just be the initial setup activity.

Option C is incorrect since this is used to run administrative tasks on on-premise SQL databases.  
<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-sync-data>

### QUESTION 174

You need to migrate data from an Azure Blob storage account to an Azure SQL Data warehouse. Which of the following actions do you need to implement for this requirement? (Choose 4)

- A. Provision an Azure SQL Data Warehouse instance
- B. Connect to the Blob storage container via SQL Server Management Studio
- C. Create an Azure Blob storage container
- D. Run the T-SQL statements to load the data
- E. Connect to the Azure SQL Data warehouse via SQL Server Management Studio
- F. Build external tables by using Azure portal
- G. Build external tables by using SQL Server Management Studio

**Answer:** ADEG

#### **Explanation:**

You first need to create an Azure SQL Data Warehouse instance.

Then you need to connect to the data warehouse via SQL Server Management Studio

Then create external tables to the Azure Blob storage account.

And then finally use T-SQL statements to load the data.

This is also given as an example in github as part of the Microsoft documentation on loading data from Azure Blob to an Azure SQL data warehouse.

This tutorial uses PolyBase to load New York Taxicab data from a global Azure blob to Azure SQL Data Warehouse. The tutorial uses the [Azure portal](#) and [SQL Server Management Studio \(SSMS\)](#) to:

[!div class="checkboxlist"]

- Create a data warehouse in the Azure portal
- Set up a server-level firewall rule in the Azure portal
- Connect to the data warehouse with SSMS
- Create a user designated for loading data
- Create external tables for data in Azure blob storage
- Use the CTAS T-SQL statement to load data into your data warehouse
- View the progress of data as it is loading
- Create statistics on the newly loaded data

Option B is incorrect because you can't connect to Blob storage from SQL Server Management Studio.

Option C is incorrect because you already have the blob data in place.

Option F is incorrect because you need to build the external tables in SQL Server Management Studio.

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/sql-data-warehouse/load-data-from-azure-blob-storage-using-polybase.md>