



# Microsoft

## Exam Questions DP-200

Implementing an Azure Data Solution

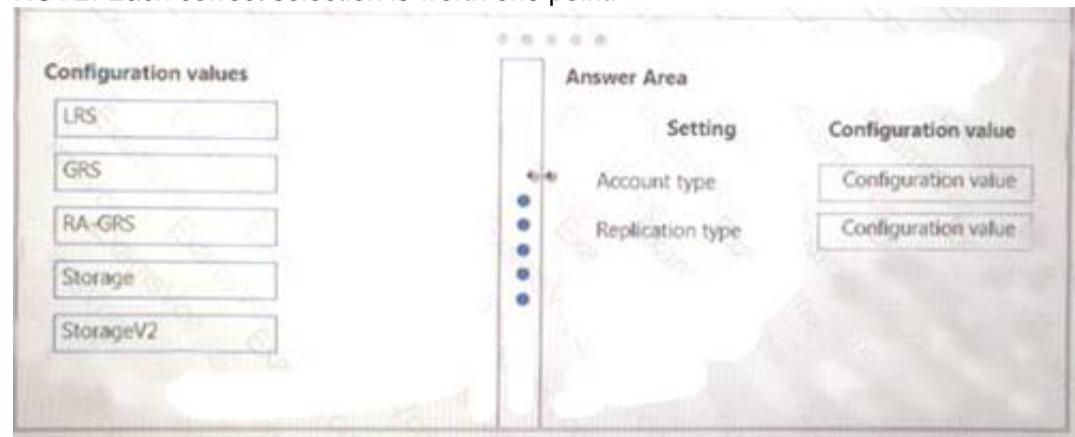
NEW QUESTION 1

- (Exam Topic 1)

You need to provision the polling data storage account.

How should you configure the storage account? To answer, drag the appropriate Configuration Value to the correct Setting. Each Configuration 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.

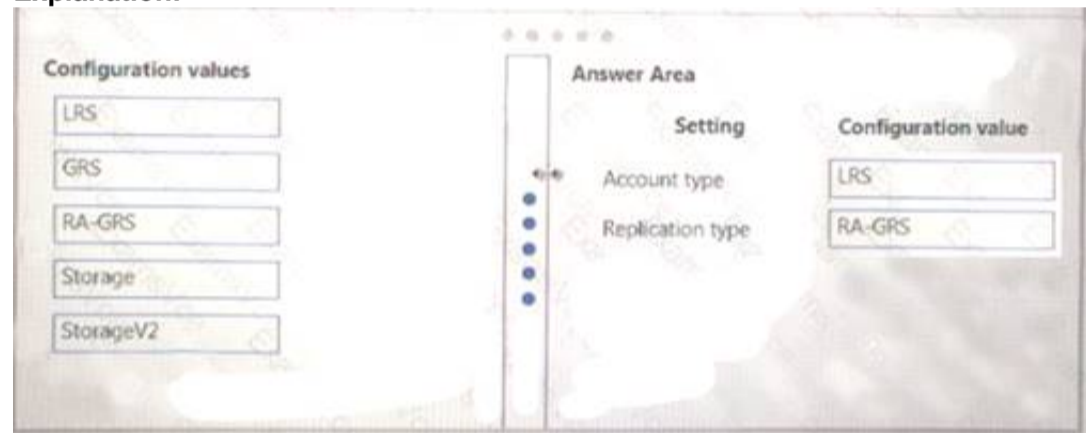
NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 2

- (Exam Topic 1)

You need to ensure polling data security requirements are met.

Which security technologies should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Context	Security technology
SQL Server	Azure Active Directory user
	Domain Active Directory user
	Managed Identity
PolyBase	Database scoped credential
	Database encryption key
	Application role

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Azure Active Directory user Scenario:

Access to polling data must set on a per-active directory user basis

Box 2: DataBase Scoped Credential

SQL Server uses a database scoped credential to access non-public Azure blob storage or Kerberos-secured Hadoop clusters with PolyBase.

PolyBase cannot authenticate by using Azure AD authentication. References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-scoped-credential-transact-sql>

NEW QUESTION 3

- (Exam Topic 2)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some questions sets might have more than one correct solution, while others might not have a correct solution.

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

You need to configure data encryption for external applications. Solution:

1. Access the Always Encrypted Wizard in SQL Server Management Studio
2. Select the column to be encrypted
3. Set the encryption type to Randomized
4. Configure the master key to use the Windows Certificate Store
5. Validate configuration results and deploy the solution Does the solution meet the goal?

- A. Yes  
B. No

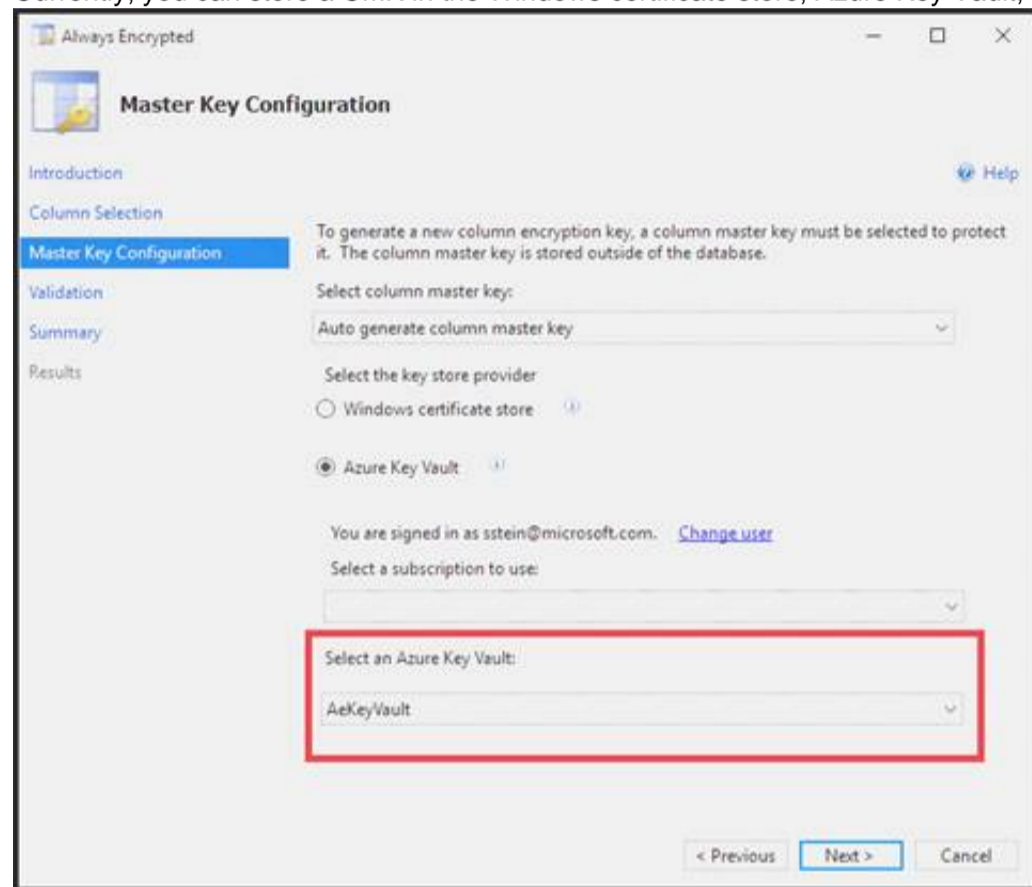
**Answer: B**

**Explanation:**

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

Note: The Master Key Configuration page is where you set up your CMK (Column Master Key) and select the key store provider where the CMK will be stored.

Currently, you can store a CMK in the Windows certificate store, Azure Key Vault, or a hardware security module (HSM).



References:

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

**NEW QUESTION 4**

- (Exam Topic 3)

Your company manages on-premises Microsoft SQL Server pipelines by using a custom solution.

The data engineering team must implement a process to pull data from SQL Server and migrate it to Azure Blob storage. The process must orchestrate and manage the data lifecycle.

You need to configure Azure Data Factory to connect to the on-premises SQL Server database.

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.

Actions	Answer Area
Create an Azure Data Factory resource.	
Configure a self-hosted integration runtime.	
Create a virtual private network (VPN) connection from on-premises to Microsoft Azure.	
Create a database master key on SQL Server.	
Backup the database and send it Azure Blob storage.	
Configure the on-premises SQL Server instance with an integration runtime.	

- A. Mastered  
B. Not Mastered

**Answer: A**

**Explanation:**

Step 1: Create a virtual private network (VPN) connection from on-premises to Microsoft Azure.

You can also use IPsec VPN or Azure ExpressRoute to further secure the communication channel between your on-premises network and Azure.

Azure Virtual Network is a logical representation of your network in the cloud. You can connect an on-premises network to your virtual network by setting up IPsec VPN (site-to-site) or ExpressRoute (private peering).

Step 2: Create an Azure Data Factory resource. Step 3: Configure a self-hosted integration runtime.

You create a self-hosted integration runtime and associate it with an on-premises machine with the SQL Server database. The self-hosted integration runtime is the component that copies data from the SQL Server database on your machine to Azure Blob storage.

Note: A self-hosted integration runtime can run copy activities between a cloud data store and a data store in a private network, and it can dispatch transform activities against compute resources in an on-premises network or an Azure virtual network. The installation of a self-hosted integration runtime needs on an on-premises machine or a virtual machine (VM) inside a private network.

References:

<https://docs.microsoft.com/en-us/azure/data-factory/tutorial-hybrid-copy-powershell>

### NEW QUESTION 5

- (Exam Topic 3)

Each day, company plans to store hundreds of files in Azure Blob Storage and Azure Data Lake Storage. The company uses the parquet format.

You must develop a pipeline that meets the following requirements:

- ▶ Process data every six hours
- ▶ Offer interactive data analysis capabilities
- ▶ Offer the ability to process data using solid-state drive (SSD) caching
- ▶ Use Directed Acyclic Graph(DAG) processing mechanisms
- ▶ Provide support for REST API calls to monitor processes
- ▶ Provide native support for Python
- ▶ Integrate with Microsoft Power BI

You need to select the appropriate data technology to implement the pipeline. Which data technology should you implement?

- A. Azure SQL Data Warehouse
- B. HDInsight Apache Storm cluster
- C. Azure Stream Analytics
- D. HDInsight Apache Hadoop cluster using MapReduce
- E. HDInsight Spark cluster

**Answer: B**

#### Explanation:

Storm runs topologies instead of the Apache Hadoop MapReduce jobs that you might be familiar with. Storm topologies are composed of multiple components that are arranged in a directed acyclic graph (DAG). Data flows between the components in the graph. Each component consumes one or more data streams, and can optionally emit one or more streams.

Python can be used to develop Storm components. References:

<https://docs.microsoft.com/en-us/azure/hdinsight/storm/apache-storm-overview>

### NEW QUESTION 6

- (Exam Topic 3)

You are a data engineer implementing a lambda architecture on Microsoft Azure. You use an open-source big data solution to collect, process, and maintain data. The analytical data store performs poorly.

You must implement a solution that meets the following requirements:

- ▶ Provide data warehousing
- ▶ Reduce ongoing management activities
- ▶ Deliver SQL query responses in less than one second

You need to create an HDInsight cluster to meet the requirements. Which type of cluster should you create?

- A. Interactive Query
- B. Apache Hadoop
- C. Apache HBase
- D. Apache Spark

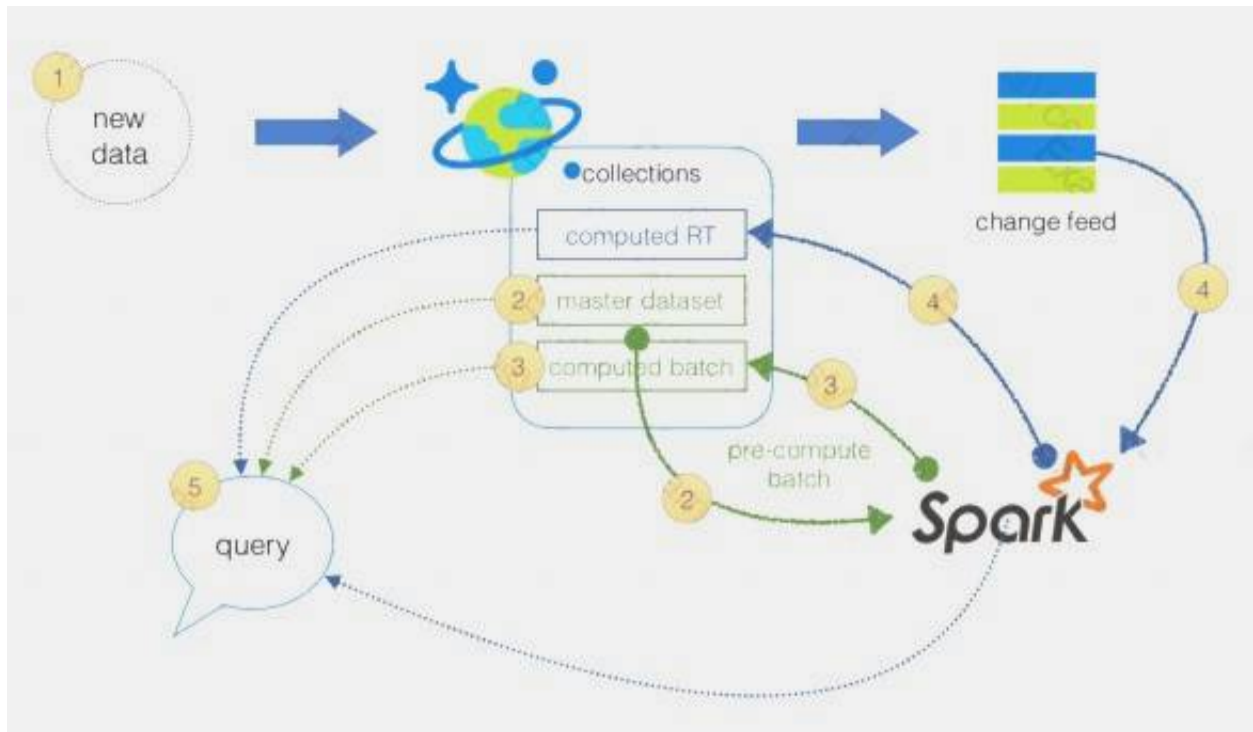
**Answer: D**

#### Explanation:

Lambda Architecture with Azure:

Azure offers you a combination of following technologies to accelerate real-time big data analytics:

- ▶ Azure Cosmos DB, a globally distributed and multi-model database service.
- ▶ Apache Spark for Azure HDInsight, a processing framework that runs large-scale data analytics applications.
- ▶ The Spark to Azure Cosmos DB Connector



Note: Lambda architecture is a data-processing architecture designed to handle massive quantities of data by taking advantage of both batch processing and stream processing methods, and minimizing the latency involved in querying big data.

References:

<https://sqlwithmanoj.com/2018/02/16/what-is-lambda-architecture-and-what-azure-offers-with-its-new-cosmos->

### NEW QUESTION 7

- (Exam Topic 3)

A company plans to develop solutions to perform batch processing of multiple sets of geospatial data. You need to implement the solutions.

Which Azure services should you use? To answer, select the appropriate configuration in the answer area. NOTE: Each correct selection is worth one point.

Scenario	Tool
Use a native client application to run interactive queries and batch processes.	<input type="checkbox"/> HDInsight Tools for Visual Studio <input type="checkbox"/> Hive View <input type="checkbox"/> HDInsight REST API <input type="checkbox"/> Azure Data Factory
Use a web browser to run interactive queries and batch processes.	<input type="checkbox"/> HDInsight Tools for Visual Studio <input type="checkbox"/> Hive View <input type="checkbox"/> HDInsight REST API <input type="checkbox"/> Azure PowerShell
Develop batch processing applications that use Azure HDInsight.	<input type="checkbox"/> HDInsight Tools for Visual Studio <input type="checkbox"/> Hive View <input type="checkbox"/> HDInsight REST API <input type="checkbox"/> NoSQL database

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Scenario	Tool
Use a native client application to run interactive queries and batch processes.	<input type="checkbox"/> HDInsight Tools for Visual Studio <input type="checkbox"/> Hive View <input checked="" type="checkbox"/> HDInsight REST API <input type="checkbox"/> Azure Data Factory
Use a web browser to run interactive queries and batch processes.	<input type="checkbox"/> HDInsight Tools for Visual Studio <input type="checkbox"/> Hive View <input type="checkbox"/> HDInsight REST API <input type="checkbox"/> Azure PowerShell
Develop batch processing applications that use Azure HDInsight.	<input type="checkbox"/> HDInsight Tools for Visual Studio <input checked="" type="checkbox"/> Hive View <input type="checkbox"/> HDInsight REST API <input type="checkbox"/> NoSQL database

### NEW QUESTION 8

- (Exam Topic 3)

A company uses Azure SQL Database to store sales transaction data. Field sales employees need an offline copy of the database that includes last year's sales on their laptops when there is no internet connection available.

You need to create the offline export copy.

Which three options can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Export to a BACPAC file by using Azure Cloud Shell, and save the file to an Azure storage account
- B. Export to a BACPAC file by using SQL Server Management Studio
- C. Save the file to an Azure storage account
- D. Export to a BACPAC file by using the Azure portal
- E. Export to a BACPAC file by using Azure PowerShell and save the file locally
- F. Export to a BACPAC file by using the SqlPackage utility

**Answer:** BCE

#### NEW QUESTION 9

- (Exam Topic 3)

A company plans to analyze a continuous flow of data from a social media platform by using Microsoft Azure Stream Analytics. The incoming data is formatted as one record per row.

You need to create the input stream.

How should you complete the REST API segment? To answer, select the appropriate configuration in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

```
{
  "properties": {
    "type": "stream",
    "serialization": {
      
      "properties": {
        "fieldDelimiter": ",",
        "encoding": "UTF8"
      }
    }
  },
  "datasource": {
    
    "properties": {
      "serviceBusNamespace": "sampleServiceBus",
      "sharedAccessPolicyName": "SampleReceiver",
      "sharedAccessPolicyKey": "<PolicyKey>",
      "eventHubName": "sampleEventHub"
    }
  },
  "compression": {
    "type": "GZip"
  }
}
```

**Answer Area**

```
{
  "properties": {
    "type": "stream",
    "serialization": {
      
      "properties": {
        "fieldDelimiter": ",",
        "encoding": "UTF8"
      }
    }
  },
  "datasource": {
    
    "properties": {
      "serviceBusNamespace": "sampleServiceBus",
      "sharedAccessPolicyName": "SampleReceiver",
      "sharedAccessPolicyKey": "<PolicyKey>",
      "eventHubName": "sampleEventHub"
    }
  }
}
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Answer Area

```
{
  "properties":{
    "type":"stream",
    "serialization":{
      "type":"CSV",
      "type":"Avro",
      "type":"JSON",
    },
    "properties":{
      "fieldDelimiter":",",
      "encoding":"UTF8"
    }
  },
  "datasource":{
    "type":"Microsoft.Storage/Blob",
    "type":"Microsoft.ServiceBus/EventHub",
    "type":"Microsoft.Devices/IotHubs",
    "properties":{
      "serviceBusNamespace":"sampleServiceBus",
      "sharedAccessPolicyName":"SampleReceiver",
      "sharedAccessPolicyKey":"<PolicyKey>"
      "eventHubName":"sampleEventHub"
    }
  }
}
```

NEW QUESTION 10

- (Exam Topic 3)

You develop data engineering solutions for a company. You must migrate data from Microsoft Azure Blob storage to an Azure SQL Data Warehouse for further transformation. You need to implement the solution. 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.

Actions	Answer Area
Provision an Azure SQL Data Warehouse instance.	
Connect to the Blob storage container by using SQL Server Management Studio.	
Provision an Azure Blob storage container.	
Run Transact-SQL statements to load data.	
Connect to the Azure SQL Data Warehouse by using SQL Server Management Studio.	
Build external tables by using Azure portal.	
Build external tables by using SQL Server Management Studio.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Provision an Azure SQL Data Warehouse instance. Create a data warehouse in the Azure portal.  
Step 2: Connect to the Azure SQL Data warehouse by using SQL Server Management Studio Connect to the data warehouse with SSMS (SQL Server Management Studio)  
Step 3: Build external tables by using the SQL Server Management Studio  
Create external tables for data in Azure blob storage.  
You are ready to begin the process of loading data into your new data warehouse. You use external tables to load data from the Azure storage blob.  
Step 4: Run Transact-SQL statements to load data.  
You can use the CREATE TABLE AS SELECT (CTAS) T-SQL statement to load the data from Azure Storage Blob into new tables in your data warehouse.  
References:  
<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/sql-data-warehouse/load-data-from-azure-blo>

NEW QUESTION 10

- (Exam Topic 3)

You are developing a data engineering solution for a company. The solution will store a large set of key-value pair data by using Microsoft Azure Cosmos DB The solution has the following requirements:

- Data must be partitioned into multiple containers.
- Data containers must be configured separately.
- Data must be accessible from applications hosted around the world.
- The solution must minimize latency. You need to provision Azure Cosmos DB

- A. Configure account-level throughput.
- B. Provision an Azure Cosmos DB account with the Azure Table API Enable geo-redundancy.
- C. Configure table-level throughput
- D. Replicate the data globally by manually adding regions lo the Azure Cosmos DB account.
- E. Provision an Azure Cosmos DB account with the Azure Table AP

F. Enable multi-region writes.

**Answer:** A

#### NEW QUESTION 11

- (Exam Topic 3)

The data engineering team manages Azure HDInsight clusters. The team spends a large amount of time creating and destroying clusters daily because most of the data pipeline process runs in minutes.

You need to implement a solution that deploys multiple HDInsight clusters with minimal effort. What should you implement?

- A. Azure Databricks
- B. Azure Traffic Manager
- C. Azure Resource Manager templates
- D. Ambari web user interface

**Answer:** C

#### Explanation:

A Resource Manager template makes it easy to create the following resources for your application in a single, coordinated operation:

- ▶ HDInsight clusters and their dependent resources (such as the default storage account).
- ▶ Other resources (such as Azure SQL Database to use Apache Sqoop).

In the template, you define the resources that are needed for the application. You also specify deployment parameters to input values for different environments. The template consists of JSON and expressions that you use to construct values for your deployment.

References:

<https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hadoop-create-linux-clusters-arm-templates>

#### NEW QUESTION 12

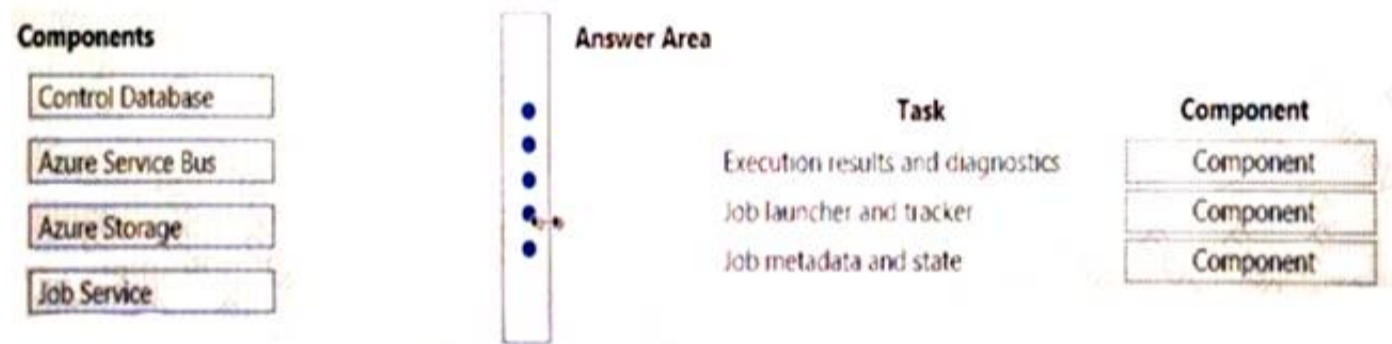
- (Exam Topic 3)

Your company uses Microsoft Azure SQL Database configured with Elastic pool. You use Elastic Database jobs to run queries across all databases in the pool.

You need to analyze, troubleshoot, and report on components responsible for running Elastic Database jobs. You need to determine the component responsible for running job service tasks.

Which components should you use for each Elastic pool job service task? To answer, drag the appropriate component to the correct task. Each component may be used once, more than once, or not at all. You may

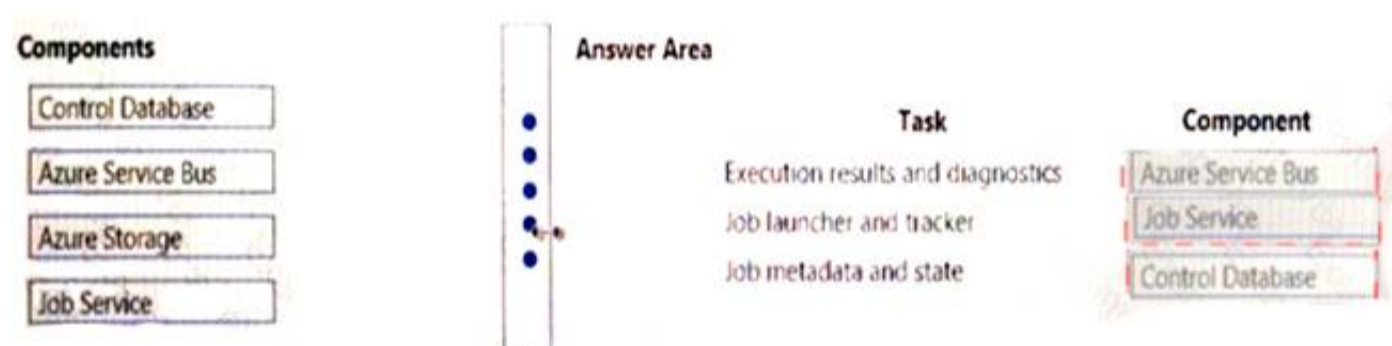
need to drag the split bar between panes or scroll to view content. NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:



#### NEW QUESTION 13

- (Exam Topic 3)

A company builds an application to allow developers to share and compare code. The conversations, code snippets, and links shared by people in the application are stored in a Microsoft Azure SQL Database instance. The application allows for searches of historical conversations and code snippets.

When users share code snippets, the code snippet is compared against previously shared code snippets by using a combination of Transact-SQL functions including SUBSTRING, FIRST\_VALUE, and SQRT. If a match is found, a link to the match is added to the conversation.

Customers report the following issues:

- ▶ Delays occur during live conversations
- ▶ A delay occurs before matching links appear after code snippets are added to conversations

You need to resolve the performance issues.

Which technologies should you use? To answer, drag the appropriate technologies to the correct issues. Each technology 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.

Technologies	Answer Area	
	Issue	Technology
columnstore index	Delays in conversations	
non-durable table		
materialized view	Delays in match links	
memory-optimized table		

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: memory-optimized table

In-Memory OLTP can provide great performance benefits for transaction processing, data ingestion, and transient data scenarios.

Box 2: materialized view

To support efficient querying, a common solution is to generate, in advance, a view that materializes the data in a format suited to the required results set. The Materialized View pattern describes generating prepopulated views of data in environments where the source data isn't in a suitable format for querying, where generating a suitable query is difficult, or where query performance is poor due to the nature of the data or the data store.

These materialized views, which only contain data required by a query, allow applications to quickly obtain the information they need. In addition to joining tables or combining data entities, materialized views can include the current values of calculated columns or data items, the results of combining values or executing transformations on the data items, and values specified as part of the query. A materialized view can even be optimized for just a single query.

References:

<https://docs.microsoft.com/en-us/azure/architecture/patterns/materialized-view>

**NEW QUESTION 15**

- (Exam Topic 3)

You configure monitoring for a Microsoft Azure SQL Data Warehouse implementation. The implementation uses PolyBase to load data from comma-separated value (CSV) files stored in Azure Data Lake Gen 2 using an external table.

Files with an invalid schema cause errors to occur. You need to monitor for an invalid schema error. For which error should you monitor?

- A. EXTERNAL TABLE access failed due to internal error: 'Java exception raised on call to HdfsBridge\_Connect: Error[com.microsoft.polybase.client.KerberosSecureLogin] occurred while accessing external files.'
- B. EXTERNAL TABLE access failed due to internal error: 'Java exception raised on call to HdfsBridge\_Connect: Error [No FileSystem for scheme: wasbs] occurred while accessing external file.'
- C. Cannot execute the query "Remote Query" against OLE DB provider "SQLNCLI11": for linked server "(null)", Query aborted- the maximum reject threshold (orows) was reached while regarding from an external source: 1 rows rejected out of total 1 rows processed.
- D. EXTERNAL TABLE access failed due to internal error: 'Java exception raised on call to HdfsBridge\_Connect: Error [Unable to instantiate LoginClass] occurredwhile accessing external files.'

**Answer:** C

**Explanation:**

Customer Scenario:

SQL Server 2016 or SQL DW connected to Azure blob storage. The CREATE EXTERNAL TABLE DDL points to a directory (and not a specific file) and the directory contains files with different schemas.

SSMS Error:

Select query on the external table gives the following error: Msg 7320, Level 16, State 110, Line 14

Cannot execute the query "Remote Query" against OLE DB provider "SQLNCLI11" for linked server "(null)". Query aborted-- the maximum reject threshold (0 rows) was reached while reading from an external source: 1 rows rejected out of total 1 rows processed.

Possible Reason:

The reason this error happens is because each file has different schema. The PolyBase external table DDL when pointed to a directory recursively reads all the files in that directory. When a column or data type mismatch happens, this error could be seen in SSMS.

Possible Solution:

If the data for each table consists of one file, then use the filename in the LOCATION section prepended by the directory of the external files. If there are multiple files per table, put each set of files into different directories in Azure Blob Storage and then you can point LOCATION to the directory instead of a particular file.

The latter suggestion is the best practices recommended by SQLCAT even if you have one file per table.

**NEW QUESTION 19**

- (Exam Topic 3)

Note: This question is part of series of questions that present the same scenario. Each question in the series contain a unique solution. Determine whether the solution meets the stated goals.

You develop a data ingestion process that will import data to a Microsoft Azure SQL Data Warehouse. The data to be ingested resides in parquet files stored in an Azure Data Lake Gen 2 storage account.

You need to load the data from the Azure Data Lake Gen 2 storage account into the Azure SQL Data Warehouse.

Solution:

1. Use Azure Data Factory to convert the parquet files to CSV files
2. Create an external data source pointing to the Azure storage account
3. Create an external file format and external table using the external data source
4. Load the data using the INSERT...SELECT statement Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

There is no need to convert the parquet files to CSV files.  
You load the data using the CREATE TABLE AS SELECT statement. References:  
<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-load-from-azure-data-lake-store>

NEW QUESTION 21

- (Exam Topic 3)  
Your company plans to create an event processing engine to handle streaming data from Twitter. The data engineering team uses Azure Event Hubs to ingest the streaming data.  
You need to implement a solution that uses Azure Databricks to receive the streaming data from the Azure Event Hubs.  
Which three actions should you recommend be performed in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

Create and configure a Notebook that consumes the streaming data.

Import data from Blob storage.

Use Environment variables to define the Apache Spark connection.

Configure an ODBC or JDBC Connector.

Deploy the Azure Databricks service.

Deploy a Spark cluster and then attach the required libraries to the cluster.

>

<

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

Answer Area

Create and configure a Notebook that consumes the streaming data.

Import data from Blob storage.

Use Environment variables to define the Apache Spark connection.

Configure an ODBC or JDBC Connector.

Deploy the Azure Databricks service.

Deploy a Spark cluster and then attach the required libraries to the cluster.

Import data from Blob storage.

Configure an ODBC or JDBC Connector.

Use Environment variables to define the Apache S connection.

NEW QUESTION 26

- (Exam Topic 3)  
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 toad the data into a large table called FactSalesOrderDetails.  
You need to configure Azure SQL Data Warehouse 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.

Actions

Create an external file format to map the parquet files.

Load the data to a staging table

Create the external table FactSalesOrderDetails.

Enable Transparent Data Encryption.

Create an external data source for Azure Blob storage.

Create a master key on database

Configure PolyBase to use Azure Blob storage.

Answer Area

>

<

- A. Mastered  
 B. Not Mastered

**Answer: A**

**Explanation:**

Actions

Create an external file format to map the parquet files.

Load the data to a staging table

Create the external table FactSalesOrderDetails.

Enable Transparent Data Encryption.

Create an external data source for Azure Blob storage.

Create a master key on database

Configure PolyBase to use Azure Blob storage.

Answer Area

Enable Transparent Data Encryption.

Configure PolyBase to use Azure Blob storage.

Load the data to a staging table

Create an external file format to map the parquet files.

#### NEW QUESTION 27

- (Exam Topic 3)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You develop a data ingestion process that will import data to a Microsoft Azure SQL Data Warehouse. The data to be ingested resides in parquet files stored in an Azure Data lake Gen 2 storage account.

You need to load the data from the Azure Data Lake Gen 2 storage account into the Azure SQL Data Warehouse.

Solution:

1. Create an external data source pointing to the Azure storage account
2. Create a workload group using the Azure storage account name as the pool name
3. Load the data using the INSERT...SELECT statement

Does the solution meet the goal?

- A. Yes  
 B. No

**Answer: B**

**Explanation:**

You need to create an external file format and external table using the external data source. You then load the data using the CREATE TABLE AS SELECT statement.

References:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-load-from-azure-data-lake-store>

#### NEW QUESTION 29

- (Exam Topic 3)

You manage a financial computation data analysis process. Microsoft Azure virtual machines (VMs) run the process in daily jobs, and store the results in virtual hard drives (VHDs.)

The VMs product results using data from the previous day and store the results in a snapshot of the VHD. When a new month begins, a process creates a new VHD.

You must implement the following data retention requirements:

- Daily results must be kept for 90 days
- Data for the current year must be available for weekly reports
- Data from the previous 10 years must be stored for auditing purposes
- Data required for an audit must be produced within 10 days of a request. You need to enforce the data retention requirements while minimizing cost.

How should you configure the lifecycle policy? To answer, drag the appropriate JSON segments to the correct locations. Each JSON segment may be used once,

more than once, or not at all. You may need to drag the split bat between panes or scroll to view content.  
NOTE: Each correct selection is worth one point.

Code segments	Answer Area
<div>delete</div>	<pre>{   "version": "0.5",   "rules": [     {       "name": "dataRetention",       "type": "Lifecycle",       "definition": {         "actions": {           "": {             "": {               "daysAfterModificationGreaterThan": 365             },             "": {               "daysAfterModificationGreaterThan": 3650             }           }         },         "": {           "": {             "daysAfterCreationGreaterThan": 90           }         }       }     }   ] }</pre>
<div>blockBlob</div>	
<div>baseBlob</div>	
<div>snapshot</div>	
<div>tierToCool</div>	
<div>tierToArchive</div>	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

The Set-AzStorageAccountManagementPolicy cmdlet creates or modifies the management policy of an Azure Storage account.  
Example: Create or update the management policy of a Storage account with ManagementPolicy rule objects.  
Action -BaseBlobAction Delete -daysAfterModificationGreaterThan 100  
PS C:\>\$action1 = Add-AzStorageAccountManagementPolicyAction -InputObject \$action1 -BaseBlobAction TierToArchive -daysAfterModificationGreaterThan 50  
PS C:\>\$action1 = Add-AzStorageAccountManagementPolicyAction -InputObject \$action1 -BaseBlobAction TierToCool -daysAfterModificationGreaterThan 30  
PS C:\>\$action1 = Add-AzStorageAccountManagementPolicyAction -InputObject \$action1 -SnapshotAction Delete -daysAfterCreationGreaterThan 100  
PS C:\>\$filter1 = New-AzStorageAccountManagementPolicyFilter -PrefixMatch ab,cd  
PS C:\>\$rule1 = New-AzStorageAccountManagementPolicyRule -Name Test -Action \$action1 -Filter \$filter1  
PS C:\>\$action2 = Add-AzStorageAccountManagementPolicyAction -BaseBlobAction Delete -daysAfterModificationGreaterThan 100  
PS C:\>\$filter2 = New-AzStorageAccountManagementPolicyFilter  
References:  
<https://docs.microsoft.com/en-us/powershell/module/az.storage/set-azstorageaccountmanagementpolicy>

NEW QUESTION 31

.....

## About ExamBible

*[Your Partner of IT Exam](#)*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

- (Exam Topic 1)

You need to provision the polling data storage account.

How should you configure the storage account? To answer, drag the appropriate Configuration Value to the correct Setting. Each Configuration 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.

Configuration values

LRS

GRS

RA-GRS

Storage

StorageV2

Answer Area

Setting

Account type

Replication type

Configuration value

Configuration value

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Configuration values

LRS

GRS

RA-GRS

Storage

StorageV2

Answer Area

Setting

Account type

Replication type

Configuration value

LRS

RA-GRS

NEW QUESTION 2

- (Exam Topic 1)

You need to ensure polling data security requirements are met.

Which security technologies should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Context	Security technology	
SQL Server	Azure Active Directory user	✓
	Domain Active Directory user	
	Managed Identity	
PolyBase	Database scoped credential	✓
	Database encryption key	
	Application role	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Azure Active Directory user Scenario:

Access to polling data must set on a per-active directory user basis

Box 2: DataBase Scoped Credential

SQL Server uses a database scoped credential to access non-public Azure blob storage or Kerberos-secured Hadoop clusters with PolyBase.

PolyBase cannot authenticate by using Azure AD authentication. References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-scoped-credential-transact-sql>

NEW QUESTION 3

- (Exam Topic 2)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some questions sets might have more than one correct solution, while others might not have a correct solution.

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

You need to configure data encryption for external applications. Solution:

1. Access the Always Encrypted Wizard in SQL Server Management Studio
2. Select the column to be encrypted
3. Set the encryption type to Randomized
4. Configure the master key to use the Windows Certificate Store
5. Validate configuration results and deploy the solution Does the solution meet the goal?

- A. Yes  
B. No

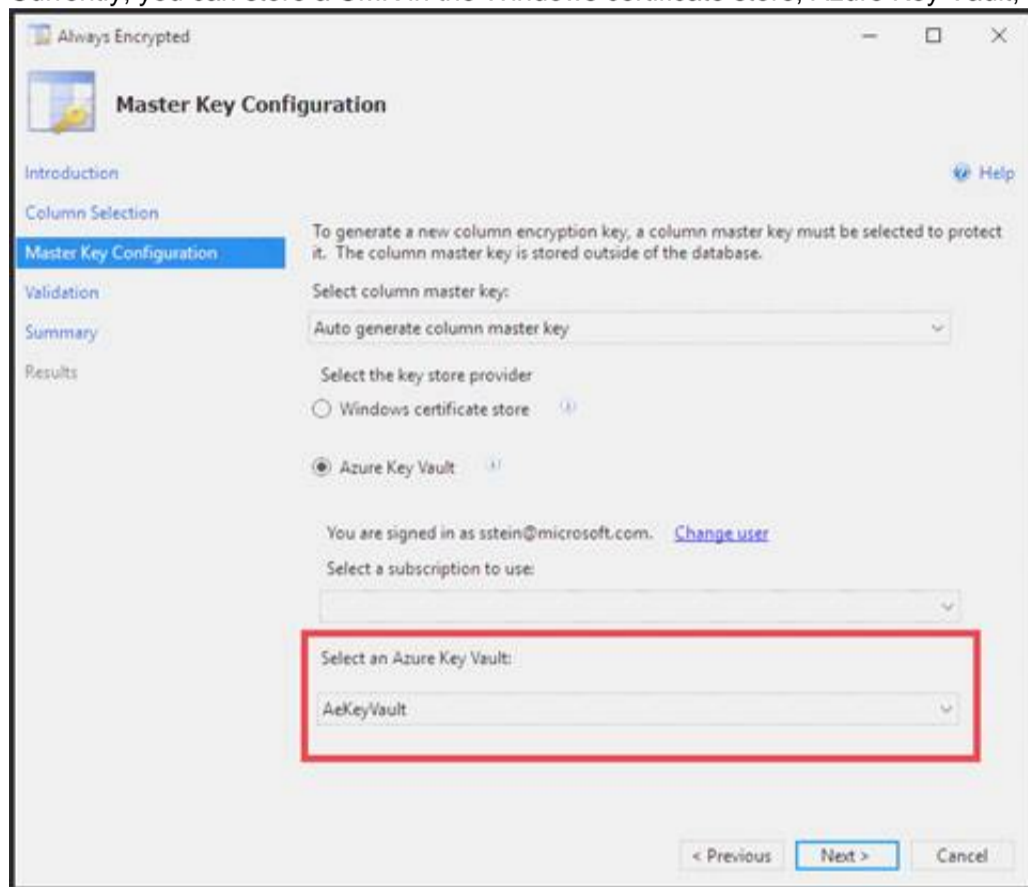
**Answer: B**

**Explanation:**

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

Note: The Master Key Configuration page is where you set up your CMK (Column Master Key) and select the key store provider where the CMK will be stored.

Currently, you can store a CMK in the Windows certificate store, Azure Key Vault, or a hardware security module (HSM).



References:

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

**NEW QUESTION 4**

- (Exam Topic 3)

Your company manages on-premises Microsoft SQL Server pipelines by using a custom solution.

The data engineering team must implement a process to pull data from SQL Server and migrate it to Azure Blob storage. The process must orchestrate and manage the data lifecycle.

You need to configure Azure Data Factory to connect to the on-premises SQL Server database.

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.

Actions	Answer Area
Create an Azure Data Factory resource.	
Configure a self-hosted integration runtime.	
Create a virtual private network (VPN) connection from on-premises to Microsoft Azure.	
Create a database master key on SQL Server.	
Backup the database and send it Azure Blob storage.	
Configure the on-premises SQL Server instance with an integration runtime.	

- A. Mastered  
B. Not Mastered

**Answer: A**

**Explanation:**

Step 1: Create a virtual private network (VPN) connection from on-premises to Microsoft Azure.

You can also use IPsec VPN or Azure ExpressRoute to further secure the communication channel between your on-premises network and Azure.

Azure Virtual Network is a logical representation of your network in the cloud. You can connect an on-premises network to your virtual network by setting up IPsec VPN (site-to-site) or ExpressRoute (private peering).

Step 2: Create an Azure Data Factory resource. Step 3: Configure a self-hosted integration runtime.

You create a self-hosted integration runtime and associate it with an on-premises machine with the SQL Server database. The self-hosted integration runtime is the component that copies data from the SQL Server database on your machine to Azure Blob storage.

Note: A self-hosted integration runtime can run copy activities between a cloud data store and a data store in a private network, and it can dispatch transform activities against compute resources in an on-premises network or an Azure virtual network. The installation of a self-hosted integration runtime needs on an on-premises machine or a virtual machine (VM) inside a private network.

References:

<https://docs.microsoft.com/en-us/azure/data-factory/tutorial-hybrid-copy-powershell>

### NEW QUESTION 5

- (Exam Topic 3)

Each day, company plans to store hundreds of files in Azure Blob Storage and Azure Data Lake Storage. The company uses the parquet format.

You must develop a pipeline that meets the following requirements:

- ▶ Process data every six hours
- ▶ Offer interactive data analysis capabilities
- ▶ Offer the ability to process data using solid-state drive (SSD) caching
- ▶ Use Directed Acyclic Graph(DAG) processing mechanisms
- ▶ Provide support for REST API calls to monitor processes
- ▶ Provide native support for Python
- ▶ Integrate with Microsoft Power BI

You need to select the appropriate data technology to implement the pipeline. Which data technology should you implement?

- A. Azure SQL Data Warehouse
- B. HDInsight Apache Storm cluster
- C. Azure Stream Analytics
- D. HDInsight Apache Hadoop cluster using MapReduce
- E. HDInsight Spark cluster

**Answer: B**

#### Explanation:

Storm runs topologies instead of the Apache Hadoop MapReduce jobs that you might be familiar with. Storm topologies are composed of multiple components that are arranged in a directed acyclic graph (DAG). Data flows between the components in the graph. Each component consumes one or more data streams, and can optionally emit one or more streams.

Python can be used to develop Storm components. References:

<https://docs.microsoft.com/en-us/azure/hdinsight/storm/apache-storm-overview>

### NEW QUESTION 6

- (Exam Topic 3)

You are a data engineer implementing a lambda architecture on Microsoft Azure. You use an open-source big data solution to collect, process, and maintain data. The analytical data store performs poorly.

You must implement a solution that meets the following requirements:

- ▶ Provide data warehousing
- ▶ Reduce ongoing management activities
- ▶ Deliver SQL query responses in less than one second

You need to create an HDInsight cluster to meet the requirements. Which type of cluster should you create?

- A. Interactive Query
- B. Apache Hadoop
- C. Apache HBase
- D. Apache Spark

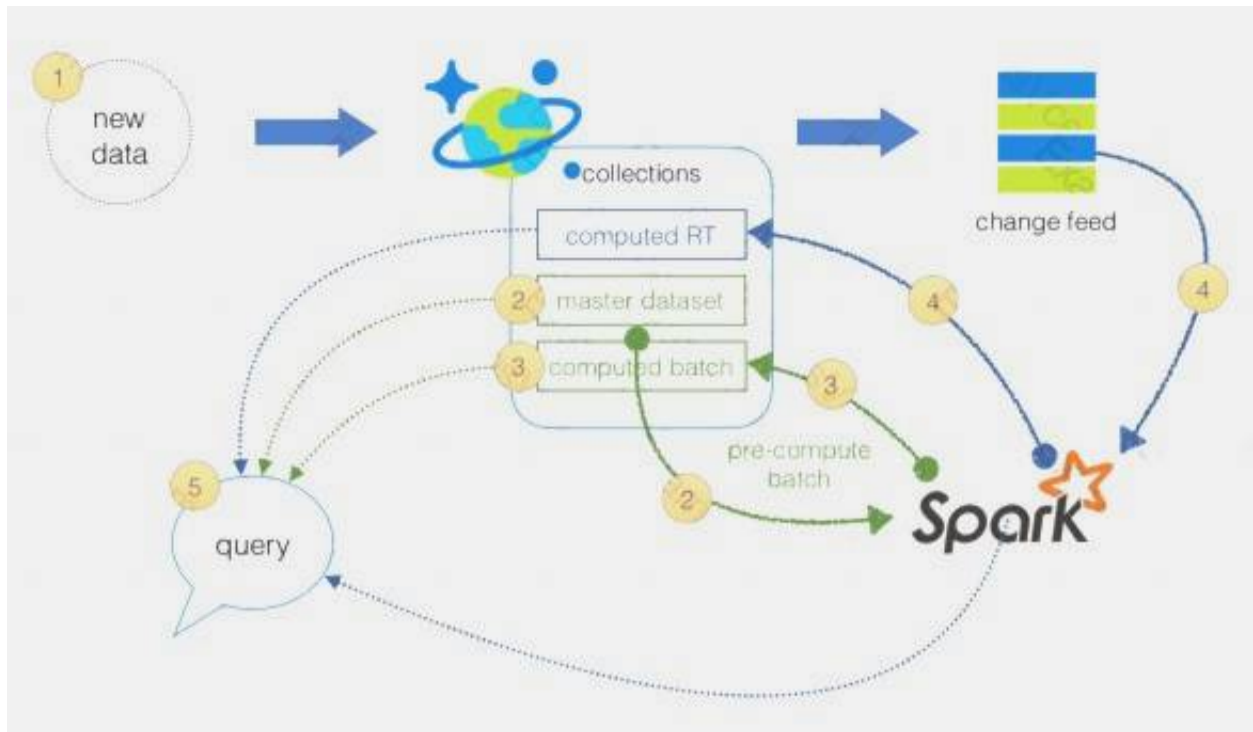
**Answer: D**

#### Explanation:

Lambda Architecture with Azure:

Azure offers you a combination of following technologies to accelerate real-time big data analytics:

- ▶ Azure Cosmos DB, a globally distributed and multi-model database service.
- ▶ Apache Spark for Azure HDInsight, a processing framework that runs large-scale data analytics applications.
- ▶ The Spark to Azure Cosmos DB Connector



Note: Lambda architecture is a data-processing architecture designed to handle massive quantities of data by taking advantage of both batch processing and stream processing methods, and minimizing the latency involved in querying big data.

References:

<https://sqlwithmanoj.com/2018/02/16/what-is-lambda-architecture-and-what-azure-offers-with-its-new-cosmos->

### NEW QUESTION 7

- (Exam Topic 3)

A company plans to develop solutions to perform batch processing of multiple sets of geospatial data. You need to implement the solutions.

Which Azure services should you use? To answer, select the appropriate configuration in the answer area. NOTE: Each correct selection is worth one point.

Scenario	Tool
Use a native client application to run interactive queries and batch processes.	<input type="checkbox"/> HDInsight Tools for Visual Studio <input type="checkbox"/> Hive View <input type="checkbox"/> HDInsight REST API <input type="checkbox"/> Azure Data Factory
Use a web browser to run interactive queries and batch processes.	<input type="checkbox"/> HDInsight Tools for Visual Studio <input type="checkbox"/> Hive View <input type="checkbox"/> HDInsight REST API <input type="checkbox"/> Azure PowerShell
Develop batch processing applications that use Azure HDInsight.	<input type="checkbox"/> HDInsight Tools for Visual Studio <input type="checkbox"/> Hive View <input type="checkbox"/> HDInsight REST API <input type="checkbox"/> NoSQL database

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Scenario	Tool
Use a native client application to run interactive queries and batch processes.	<input type="checkbox"/> HDInsight Tools for Visual Studio <input type="checkbox"/> Hive View <input checked="" type="checkbox"/> HDInsight REST API <input type="checkbox"/> Azure Data Factory
Use a web browser to run interactive queries and batch processes.	<input type="checkbox"/> HDInsight Tools for Visual Studio <input type="checkbox"/> Hive View <input type="checkbox"/> HDInsight REST API <input type="checkbox"/> Azure PowerShell
Develop batch processing applications that use Azure HDInsight.	<input type="checkbox"/> HDInsight Tools for Visual Studio <input checked="" type="checkbox"/> Hive View <input type="checkbox"/> HDInsight REST API <input type="checkbox"/> NoSQL database

### NEW QUESTION 8

- (Exam Topic 3)

A company uses Azure SQL Database to store sales transaction data. Field sales employees need an offline copy of the database that includes last year's sales on their laptops when there is no internet connection available.

You need to create the offline export copy.

Which three options can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Export to a BACPAC file by using Azure Cloud Shell, and save the file to an Azure storage account
- B. Export to a BACPAC file by using SQL Server Management Studio
- C. Save the file to an Azure storage account
- D. Export to a BACPAC file by using the Azure portal
- E. Export to a BACPAC file by using Azure PowerShell and save the file locally
- F. Export to a BACPAC file by using the SqlPackage utility

**Answer:** BCE

#### NEW QUESTION 9

- (Exam Topic 3)

A company plans to analyze a continuous flow of data from a social media platform by using Microsoft Azure Stream Analytics. The incoming data is formatted as one record per row.

You need to create the input stream.

How should you complete the REST API segment? To answer, select the appropriate configuration in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

```
{
  "properties": {
    "type": "stream",
    "serialization": {
      
      "properties": {
        "fieldDelimiter": ",",
        "encoding": "UTF8"
      }
    },
    "datasource": {
      
      "properties": {
        "serviceBusNamespace": "sampleServiceBus",
        "sharedAccessPolicyName": "SampleReceiver",
        "sharedAccessPolicyKey": "<PolicyKey>",
        "eventHubName": "sampleEventHub"
      },
      "compression": {
        "type": "GZip"
      }
    }
  }
}
```

**Answer Area**

```
{
  "properties": {
    "type": "stream",
    "serialization": {
      
      "properties": {
        "fieldDelimiter": ",",
        "encoding": "UTF8"
      }
    },
    "datasource": {
      
      "properties": {
        "serviceBusNamespace": "sampleServiceBus",
        "sharedAccessPolicyName": "SampleReceiver",
        "sharedAccessPolicyKey": "<PolicyKey>",
        "eventHubName": "sampleEventHub"
      }
    }
  }
}
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Answer Area

```
{
  "properties":{
    "type":"stream",
    "serialization":{
      "type":"CSV",
      "type":"Avro",
      "type":"JSON",
    },
    "properties":{
      "fieldDelimiter":",",
      "encoding":"UTF8"
    }
  },
  "datasource":{
    "type":"Microsoft.Storage/Blob",
    "type":"Microsoft.ServiceBus/EventHub",
    "type":"Microsoft.Devices/IotHubs",
    "properties":{
      "serviceBusNamespace":"sampleServiceBus",
      "sharedAccessPolicyName":"SampleReceiver",
      "sharedAccessPolicyKey":"<PolicyKey>"
      "eventHubName":"sampleEventHub"
    }
  }
}
```

NEW QUESTION 10

- (Exam Topic 3)

You develop data engineering solutions for a company. You must migrate data from Microsoft Azure Blob storage to an Azure SQL Data Warehouse for further transformation. You need to implement the solution. 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.

Actions	Answer Area
Provision an Azure SQL Data Warehouse instance.	
Connect to the Blob storage container by using SQL Server Management Studio.	
Provision an Azure Blob storage container.	
Run Transact-SQL statements to load data.	
Connect to the Azure SQL Data Warehouse by using SQL Server Management Studio.	
Build external tables by using Azure portal.	
Build external tables by using SQL Server Management Studio.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Provision an Azure SQL Data Warehouse instance. Create a data warehouse in the Azure portal.  
Step 2: Connect to the Azure SQL Data warehouse by using SQL Server Management Studio Connect to the data warehouse with SSMS (SQL Server Management Studio)  
Step 3: Build external tables by using the SQL Server Management Studio  
Create external tables for data in Azure blob storage.  
You are ready to begin the process of loading data into your new data warehouse. You use external tables to load data from the Azure storage blob.  
Step 4: Run Transact-SQL statements to load data.  
You can use the CREATE TABLE AS SELECT (CTAS) T-SQL statement to load the data from Azure Storage Blob into new tables in your data warehouse.  
References:  
<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/sql-data-warehouse/load-data-from-azure-blo>

NEW QUESTION 10

- (Exam Topic 3)

You are developing a data engineering solution for a company. The solution will store a large set of key-value pair data by using Microsoft Azure Cosmos DB The solution has the following requirements:

- Data must be partitioned into multiple containers.
- Data containers must be configured separately.
- Data must be accessible from applications hosted around the world.
- The solution must minimize latency. You need to provision Azure Cosmos DB

- A. Configure account-level throughput.
- B. Provision an Azure Cosmos DB account with the Azure Table API Enable geo-redundancy.
- C. Configure table-level throughput
- D. Replicate the data globally by manually adding regions lo the Azure Cosmos DB account.
- E. Provision an Azure Cosmos DB account with the Azure Table AP

F. Enable multi-region writes.

**Answer:** A

#### NEW QUESTION 11

- (Exam Topic 3)

The data engineering team manages Azure HDInsight clusters. The team spends a large amount of time creating and destroying clusters daily because most of the data pipeline process runs in minutes.

You need to implement a solution that deploys multiple HDInsight clusters with minimal effort. What should you implement?

- A. Azure Databricks
- B. Azure Traffic Manager
- C. Azure Resource Manager templates
- D. Ambari web user interface

**Answer:** C

#### Explanation:

A Resource Manager template makes it easy to create the following resources for your application in a single, coordinated operation:

- ▶ HDInsight clusters and their dependent resources (such as the default storage account).
- ▶ Other resources (such as Azure SQL Database to use Apache Sqoop).

In the template, you define the resources that are needed for the application. You also specify deployment parameters to input values for different environments. The template consists of JSON and expressions that you use to construct values for your deployment.

References:

<https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hadoop-create-linux-clusters-arm-templates>

#### NEW QUESTION 12

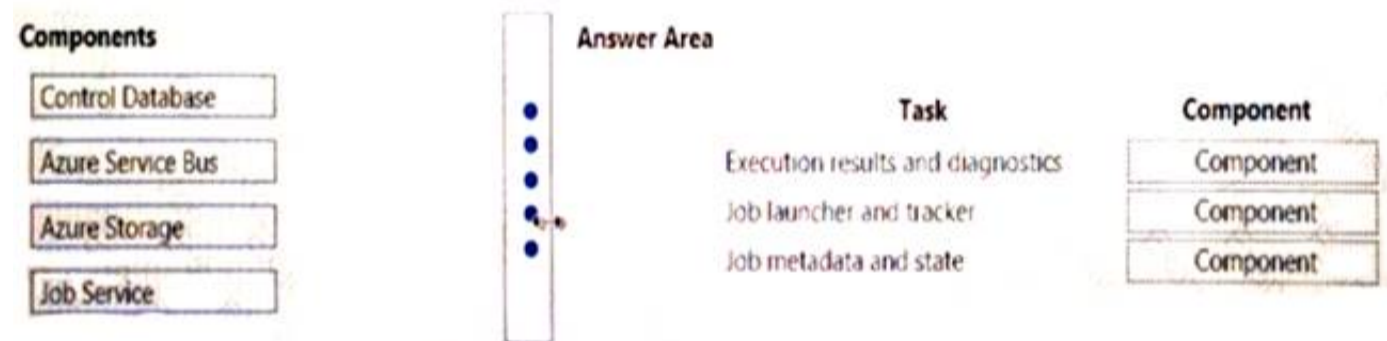
- (Exam Topic 3)

Your company uses Microsoft Azure SQL Database configured with Elastic pool. You use Elastic Database jobs to run queries across all databases in the pool.

You need to analyze, troubleshoot, and report on components responsible for running Elastic Database jobs. You need to determine the component responsible for running job service tasks.

Which components should you use for each Elastic pool job service task? To answer, drag the appropriate component to the correct task. Each component may be used once, more than once, or not at all. You may

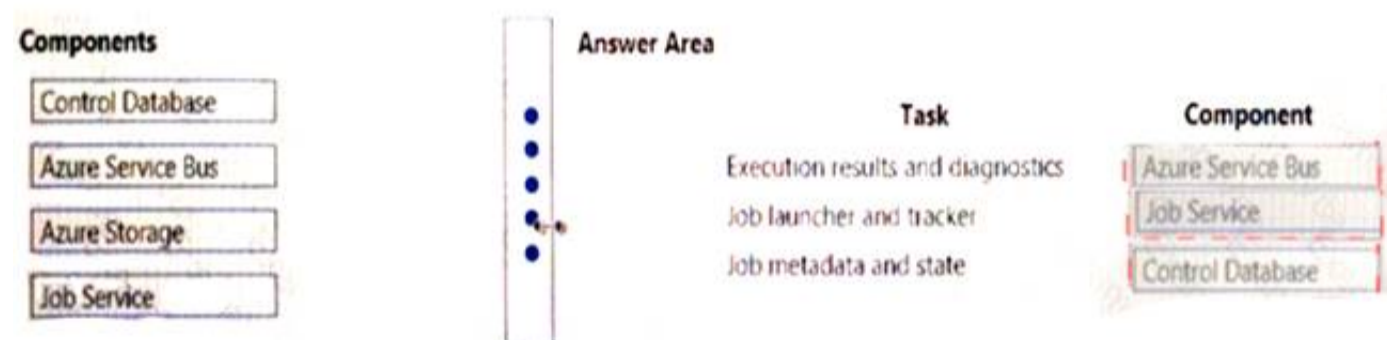
need to drag the split bar between panes or scroll to view content. NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:



#### NEW QUESTION 13

- (Exam Topic 3)

A company builds an application to allow developers to share and compare code. The conversations, code snippets, and links shared by people in the application are stored in a Microsoft Azure SQL Database instance. The application allows for searches of historical conversations and code snippets.

When users share code snippets, the code snippet is compared against previously shared code snippets by using a combination of Transact-SQL functions including SUBSTRING, FIRST\_VALUE, and SQRT. If a match is found, a link to the match is added to the conversation.

Customers report the following issues:

- ▶ Delays occur during live conversations
- ▶ A delay occurs before matching links appear after code snippets are added to conversations

You need to resolve the performance issues.

Which technologies should you use? To answer, drag the appropriate technologies to the correct issues. Each technology 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.

Technologies	Answer Area	
	Issue	Technology
columnstore index	Delays in conversations	
non-durable table	Delays in match links	
materialized view		
memory-optimized table		

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: memory-optimized table

In-Memory OLTP can provide great performance benefits for transaction processing, data ingestion, and transient data scenarios.

Box 2: materialized view

To support efficient querying, a common solution is to generate, in advance, a view that materializes the data in a format suited to the required results set. The Materialized View pattern describes generating prepopulated views of data in environments where the source data isn't in a suitable format for querying, where generating a suitable query is difficult, or where query performance is poor due to the nature of the data or the data store.

These materialized views, which only contain data required by a query, allow applications to quickly obtain the information they need. In addition to joining tables or combining data entities, materialized views can include the current values of calculated columns or data items, the results of combining values or executing transformations on the data items, and values specified as part of the query. A materialized view can even be optimized for just a single query.

References:

<https://docs.microsoft.com/en-us/azure/architecture/patterns/materialized-view>

**NEW QUESTION 15**

- (Exam Topic 3)

You configure monitoring for a Microsoft Azure SQL Data Warehouse implementation. The implementation uses PolyBase to load data from comma-separated value (CSV) files stored in Azure Data Lake Gen 2 using an external table.

Files with an invalid schema cause errors to occur. You need to monitor for an invalid schema error. For which error should you monitor?

- A. EXTERNAL TABLE access failed due to internal error: 'Java exception raised on call to HdfsBridge\_Connect: Error[com.microsoft.polybase.client.KerberosSecureLogin] occurred while accessing external files.'
- B. EXTERNAL TABLE access failed due to internal error: 'Java exception raised on call to HdfsBridge\_Connect: Error [No FileSystem for scheme: wasbs] occurred while accessing external file.'
- C. Cannot execute the query "Remote Query" against OLE DB provider "SQLNCLI11": for linked server "(null)", Query aborted- the maximum reject threshold (orows) was reached while regarding from an external source: 1 rows rejected out of total 1 rows processed.
- D. EXTERNAL TABLE access failed due to internal error: 'Java exception raised on call to HdfsBridge\_Connect: Error [Unable to instantiate LoginClass] occurredwhile accessing external files.'

**Answer: C**

**Explanation:**

Customer Scenario:

SQL Server 2016 or SQL DW connected to Azure blob storage. The CREATE EXTERNAL TABLE DDL points to a directory (and not a specific file) and the directory contains files with different schemas.

SSMS Error:

Select query on the external table gives the following error: Msg 7320, Level 16, State 110, Line 14

Cannot execute the query "Remote Query" against OLE DB provider "SQLNCLI11" for linked server "(null)". Query aborted-- the maximum reject threshold (0 rows) was reached while reading from an external source: 1 rows rejected out of total 1 rows processed.

Possible Reason:

The reason this error happens is because each file has different schema. The PolyBase external table DDL when pointed to a directory recursively reads all the files in that directory. When a column or data type mismatch happens, this error could be seen in SSMS.

Possible Solution:

If the data for each table consists of one file, then use the filename in the LOCATION section prepended by the directory of the external files. If there are multiple files per table, put each set of files into different directories in Azure Blob Storage and then you can point LOCATION to the directory instead of a particular file.

The latter suggestion is the best practices recommended by SQLCAT even if you have one file per table.

**NEW QUESTION 19**

- (Exam Topic 3)

Note: This question is part of series of questions that present the same scenario. Each question in the series contain a unique solution. Determine whether the solution meets the stated goals.

You develop a data ingestion process that will import data to a Microsoft Azure SQL Data Warehouse. The data to be ingested resides in parquet files stored in an Azure Data Lake Gen 2 storage account.

You need to load the data from the Azure Data Lake Gen 2 storage account into the Azure SQL Data Warehouse.

Solution:

1. Use Azure Data Factory to convert the parquet files to CSV files
2. Create an external data source pointing to the Azure storage account
3. Create an external file format and external table using the external data source
4. Load the data using the INSERT...SELECT statement Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

There is no need to convert the parquet files to CSV files.  
You load the data using the CREATE TABLE AS SELECT statement. References:  
<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-load-from-azure-data-lake-store>

NEW QUESTION 21

- (Exam Topic 3)  
Your company plans to create an event processing engine to handle streaming data from Twitter. The data engineering team uses Azure Event Hubs to ingest the streaming data.  
You need to implement a solution that uses Azure Databricks to receive the streaming data from the Azure Event Hubs.  
Which three actions should you recommend be performed in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

● ● ● ● ●

Actions

Create and configure a Notebook that consumes the streaming data.

Import data from Blob storage.

Use Environment variables to define the Apache Spark connection.

Configure an ODBC or JDBC Connector.

Deploy the Azure Databricks service.

Deploy a Spark cluster and then attach the required libraries to the cluster.

Answer Area

>

<

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

● ● ● ● ●

Actions

Create and configure a Notebook that consumes the streaming data.

Import data from Blob storage.

Use Environment variables to define the Apache Spark connection.

Configure an ODBC or JDBC Connector.

Deploy the Azure Databricks service.

Deploy a Spark cluster and then attach the required libraries to the cluster.

Answer Area

Import data from Blob storage.

Configure an ODBC or JDBC Connector.

Use Environment variables to define the Apache S connection.

NEW QUESTION 26

- (Exam Topic 3)  
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 toad the data into a large table called FactSalesOrderDetails.  
You need to configure Azure SQL Data Warehouse 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.

Actions

Create an external file format to map the parquet files.

Load the data to a staging table

Create the external table FactSalesOrderDetails.

Enable Transparent Data Encryption.

Create an external data source for Azure Blob storage.

Create a master key on database

Configure PolyBase to use Azure Blob storage.

Answer Area

>

<

- A. Mastered  
 B. Not Mastered

**Answer: A**

**Explanation:**

Actions

Create an external file format to map the parquet files.

Load the data to a staging table

Create the external table FactSalesOrderDetails.

Enable Transparent Data Encryption.

Create an external data source for Azure Blob storage.

Create a master key on database

Configure PolyBase to use Azure Blob storage.

Answer Area

Enable Transparent Data Encryption.

Configure PolyBase to use Azure Blob storage.

Load the data to a staging table

Create an external file format to map the parquet files.

## NEW QUESTION 27

- (Exam Topic 3)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You develop a data ingestion process that will import data to a Microsoft Azure SQL Data Warehouse. The data to be ingested resides in parquet files stored in an Azure Data lake Gen 2 storage account.

You need to load the data from the Azure Data Lake Gen 2 storage account into the Azure SQL Data Warehouse.

Solution:

1. Create an external data source pointing to the Azure storage account
2. Create a workload group using the Azure storage account name as the pool name
3. Load the data using the INSERT...SELECT statement

Does the solution meet the goal?

- A. Yes  
 B. No

**Answer: B**

**Explanation:**

You need to create an external file format and external table using the external data source. You then load the data using the CREATE TABLE AS SELECT statement.

References:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-load-from-azure-data-lake-store>

## NEW QUESTION 29

- (Exam Topic 3)

You manage a financial computation data analysis process. Microsoft Azure virtual machines (VMs) run the process in daily jobs, and store the results in virtual hard drives (VHDs.)

The VMs product results using data from the previous day and store the results in a snapshot of the VHD. When a new month begins, a process creates a new VHD.

You must implement the following data retention requirements:

- Daily results must be kept for 90 days
- Data for the current year must be available for weekly reports
- Data from the previous 10 years must be stored for auditing purposes
- Data required for an audit must be produced within 10 days of a request. You need to enforce the data retention requirements while minimizing cost.

How should you configure the lifecycle policy? To answer, drag the appropriate JSON segments to the correct locations. Each JSON segment may be used once,

more than once, or not at all. You may need to drag the split bat between panes or scroll to view content.  
NOTE: Each correct selection is worth one point.

Code segments	Answer Area
<div>delete</div>	<pre>{   "version": "0.5",   "rules": [     {       "name": "dataRetention",       "type": "Lifecycle",       "definition": {         "actions": {           "": {             "": {               "daysAfterModificationGreaterThan": 365             },             "": {               "daysAfterModificationGreaterThan": 3650             }           }         },         "": {           "": {             "daysAfterCreationGreaterThan": 90           }         }       }     }   ] }</pre>
<div>blockBlob</div>	
<div>baseBlob</div>	
<div>snapshot</div>	
<div>tierToCool</div>	
<div>tierToArchive</div>	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

The Set-AzStorageAccountManagementPolicy cmdlet creates or modifies the management policy of an Azure Storage account.  
Example: Create or update the management policy of a Storage account with ManagementPolicy rule objects.  
Action -BaseBlobAction Delete -daysAfterModificationGreaterThan 100  
PS C:\>\$action1 = Add-AzStorageAccountManagementPolicyAction -InputObject \$action1 -BaseBlobAction TierToArchive -daysAfterModificationGreaterThan 50  
PS C:\>\$action1 = Add-AzStorageAccountManagementPolicyAction -InputObject \$action1 -BaseBlobAction TierToCool -daysAfterModificationGreaterThan 30  
PS C:\>\$action1 = Add-AzStorageAccountManagementPolicyAction -InputObject \$action1 -SnapshotAction Delete -daysAfterCreationGreaterThan 100  
PS C:\>\$filter1 = New-AzStorageAccountManagementPolicyFilter -PrefixMatch ab,cd  
PS C:\>\$rule1 = New-AzStorageAccountManagementPolicyRule -Name Test -Action \$action1 -Filter \$filter1  
PS C:\>\$action2 = Add-AzStorageAccountManagementPolicyAction -BaseBlobAction Delete -daysAfterModificationGreaterThan 100  
PS C:\>\$filter2 = New-AzStorageAccountManagementPolicyFilter  
References:  
<https://docs.microsoft.com/en-us/powershell/module/az.storage/set-azstorageaccountmanagementpolicy>

NEW QUESTION 31

.....

## Relate Links

**100% Pass Your DP-200 Exam with Examible Prep Materials**

<https://www.exambible.com/DP-200-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>