DP-200 Dumps

Implementing an Azure Data Solution

https://www.certleader.com/DP-200-dumps.html

NEW QUESTION 1
- (Exam Topic 1)
You need to ensure that phone-based polling data can be analyzed in the PollingData database.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer are and arrange them in the correct order.

Actions	Answer Area
Parameterize deployment by using Azure Integration Runtime	
Configure an Azure Logic App to deploy the deployment artifact	
Configure Azure DevOps to deploy the deployment artifact	
Create a deployment artifact containing an extracted Azure Resource Manager template	
Parameterize deployment by using the Azure Resource Manager template parameter file	
Create a deployment artifact containing a SQL Server Integration Services (SSIS) package	

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Create a deployment artifact containing an extracted Azure Resource Manager template
Parameterize deployment by using the Azure Resource Manager template parameter file
Configure Azure DevOps to deploy the deployment artifact

Scenario:
All deployments must be performed by using Azure DevOps. Deployments must use templates used in multiple environments
No credentials or secrets should be used during deployments

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 1)
You need to ensure phone-based polling data upload reliability requirements are met. How should you configure monitoring? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Setting	Value
Metric	FileCount / BlobCapacity / FileCapacity
Aggregation	Avg / Sum

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: FileCapacity
FileCapacity is the amount of storage used by the storage account's File service in bytes. Box 2: Avg
The aggregation type of the FileCapacity metric is Avg.
Scenario:
All services and processes must be resilient to a regional Azure outage.
All Azure services must be monitored by using Azure Monitor. On-premises SQL Server performance must be monitored.
References:
https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported

NEW QUESTION 4
- (Exam Topic 2)
You need to set up access to Azure SQL Database for Tier 7 and Tier 8 partners.
Which three actions should you perform in sequence? To answer, move the appropriate three actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Connect to the Database and use Azure PowerShell to create a database firewall rule	
Set the Allow Azure Services to Access Server to Disabled	
In ther Azure portal, create a database firewall rule	
In the Azure portal, create a server firewall rule	
Connect to the database and use Transact-SQL to create a database firewall rule	
Set the Allow Azure Services to Access Server setting to Enabled	

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Tier 7 and 8 data access is constrained to single endpoints managed by partners for access Step 1: Set the Allow Azure Services to Access Server setting to Disabled
Set Allow access to Azure services to OFF for the most secure configuration.
By default, access through the SQL Database firewall is enabled for all Azure services, under Allow access to Azure services. Choose OFF to disable access for all Azure services.

Note: The firewall pane has an ON/OFF button that is labeled Allow access to Azure services. The ON setting allows communications from all Azure IP addresses and all Azure subnets. These Azure IPs or subnets might not be owned by you. This ON setting is probably more open than you want your SQL Database to be. The virtual network rule feature offers much finer granular control.

Step 2: In the Azure portal, create a server firewall rule Set up SQL Database server firewall rules

Server-level IP firewall rules apply to all databases within the same SQL Database server. To set up a server-level firewall rule:

▶ In Azure portal, select SQL databases from the left-hand menu, and select your database on the SQL databases page.

▶ On the Overview page, select Set server firewall. The Firewall settings page for the database server opens.

Step 3: Connect to the database and use Transact-SQL to create a database firewall rule

Database-level firewall rules can only be configured using Transact-SQL (T-SQL) statements, and only after you've configured a server-level firewall rule.

To setup a database-level firewall rule:

▶In Object Explorer, right-click the database and select New Query.

▶EXECUTE sp_set_database_firewall_rule N'Example DB Rule','0.0.0.4','0.0.0.4';

▶ On the toolbar, select Execute to create the firewall rule. References:

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-security-tutorial

NEW QUESTION 5
- (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 Deterministic
4. Configure the master key to use the Azure Key Vault
5. Validate configuration results and deploy the solution Does the solution meet the goal?
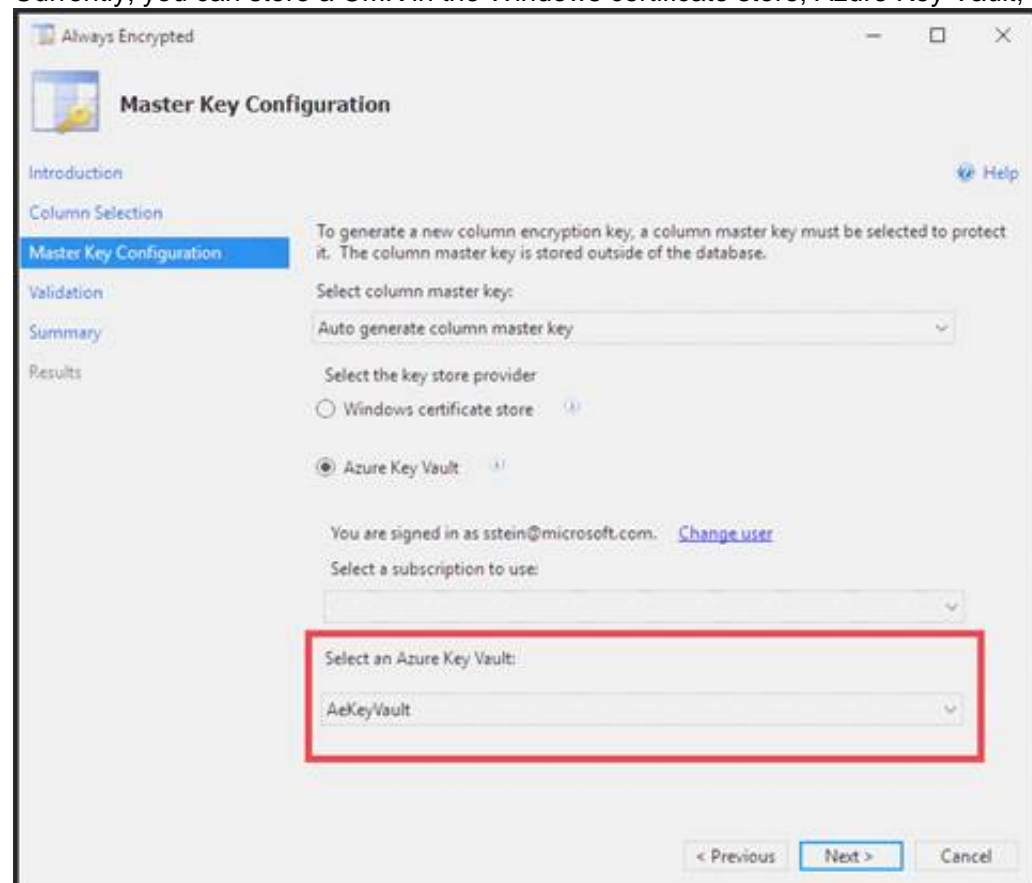
A. Yes
B. No

Answer: A

Explanation:
We 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 6
- (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 setup monitoring for tiers 6 through 8. What should you configure?

A. extended events for average storage percentage that emails data engineers
B. an alert rule to monitor CPU percentage in databases that emails data engineers
C. an alert rule to monitor CPU percentage in elastic pools that emails data engineers
D. an alert rule to monitor storage percentage in databases that emails data engineers
E. an alert rule to monitor storage percentage in elastic pools that emails data engineers

Answer: E

Explanation:
Scenario:
Tiers 6 through 8 must have unexpected resource storage usage immediately reported to data engineers.
Tier 3 and Tier 6 through Tier 8 applications must use database density on the same server and Elastic pools in a cost-effective manner.

NEW QUESTION 7
- (Exam Topic 3)
A company is planning to use Microsoft Azure Cosmos DB as the data store for an application. You have the following Azure CLI command:
az cosmosdb create -–name "cosmosdbdev1" –-resource-group "rgdev"
You need to minimize latency and expose the SQL API. How should you complete the command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

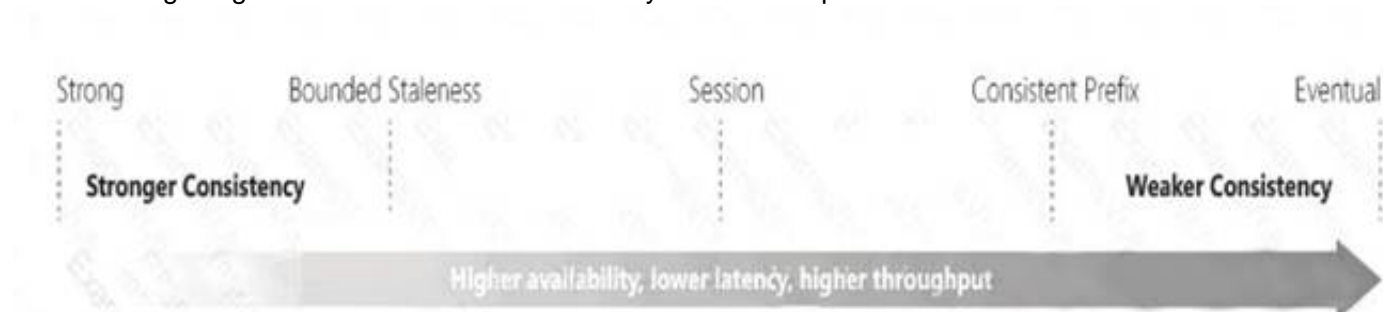| Parameter	Value
--default-consistency-level	Strong / Session / Eventual / Bounded staleness
--kind	Parse / MongoDB / GlobalDocumentDB

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Eventual
With Azure Cosmos DB, developers can choose from five well-defined consistency models on the consistency spectrum. From strongest to more relaxed, the models include strong, bounded staleness, session, consistent prefix, and eventual consistency.
The following image shows the different consistency levels as a spectrum.

Box 2: GlobalDocumentDB
Select Core(SQL) to create a document database and query by using SQL syntax.
Note: The API determines the type of account to create. Azure Cosmos DB provides five APIs: Core(SQL) and MongoDB for document databases, Gremlin for graph databases, Azure Table, and Cassandra.
References:
https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels https://docs.microsoft.com/en-us/azure/cosmos-db/create-sql-api-dotnet

NEW QUESTION 8
- (Exam Topic 3)
Contoso, Ltd. plans to configure existing applications to use Azure SQL Database. When security-related operations occur, the security team must be informed.
You need to configure Azure Monitor while minimizing administrative efforts
Which three actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Create a new action group to email alerts@contoso.com.
B. Use alerts@contoso.com as an alert email address.
C. Use all security operations as a condition.
D. Use all Azure SQL Database servers as a resource.
E. Query audit log entries as a condition.

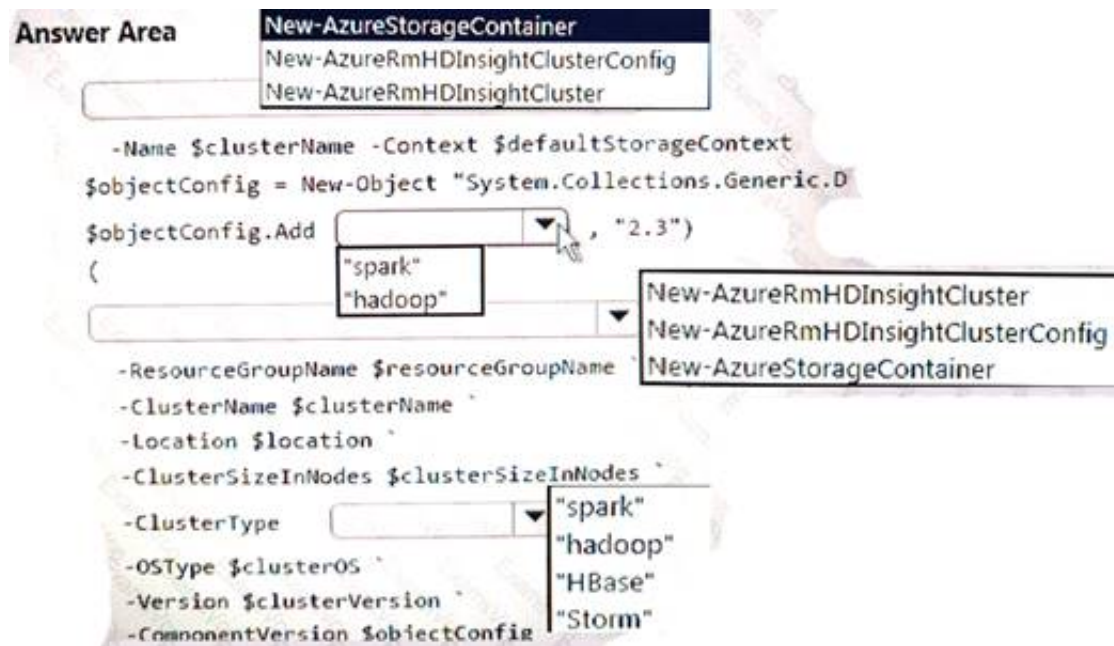**Answer:** ACE

NEW QUESTION 9
- (Exam Topic 3)
You develop data engineering solutions for a company.
A project requires an in-memory batch data processing solution.
You need to provision an HDInsight cluster for batch processing of data on Microsoft Azure.
How should you complete the PowerShell segment? To answer, select the appropriate option in the answer area.
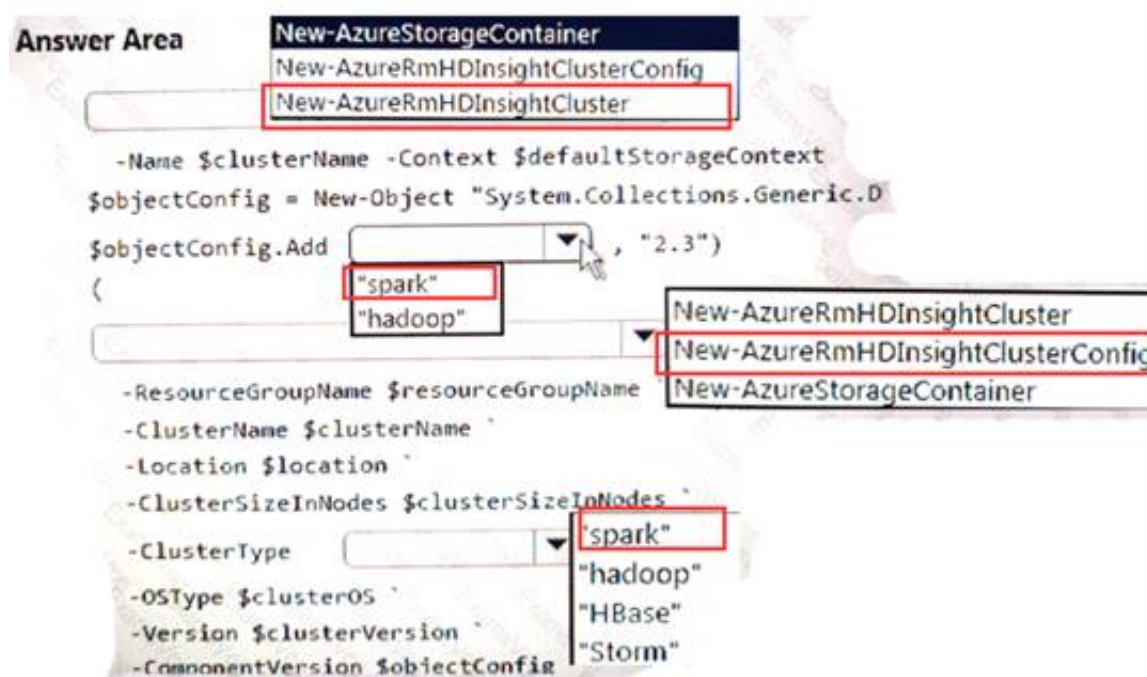NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 10
- (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 10
- (Exam Topic 3)
You manage a solution that uses Azure HDInsight clusters.
You need to implement a solution to monitor cluster performance and status. Which technology should you use?

A. Azure HDInsight .NET SDK
B. Azure HDInsight REST API
C. Ambari REST API
D. Azure Log Analytics
E. Ambari Web UI

Answer: E

Explanation:
Ambari is the recommended tool for monitoring utilization across the whole cluster. The Ambari dashboard shows easily glanceable widgets that display metrics such as CPU, network, YARN memory, and HDFS disk usage. The specific metrics shown depend on cluster type. The "Hosts" tab shows metrics for individual nodes so you can ensure the load on your cluster is evenly distributed.
The Apache Ambari project is aimed at making Hadoop management simpler by developing software for provisioning, managing, and monitoring Apache Hadoop clusters. Ambari provides an intuitive, easy-to-use Hadoop management web UI backed by its RESTful APIs.
References:
https://azure.microsoft.com/en-us/blog/monitoring-on-hdinsight-part-1-an-overview/ https://ambari.apache.org/

NEW QUESTION 13
- (Exam Topic 3)
You plan to use Microsoft Azure SQL Database instances with strict user access control. A user object must:
▶ Move with the database if it is run elsewhere
▶ Be able to create additional users
You need to create the user object with correct permissions.
Which two Transact-SQL commands should you run? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. ALTER LOGIN Mary WITH PASSWORD = 'strong_password';
B. CREATE LOGIN Mary WITH PASSWORD = 'strong_password';
C. ALTER ROLE db_owner ADD MEMBER Mary;
D. CREATE USER Mary WITH PASSWORD = 'strong_password';
E. GRANT ALTER ANY USER TO Mary;

Answer: CD

Explanation:
C: ALTER ROLE adds or removes members to or from a database role, or changes the name of a user-defined database role.
Members of the db_owner fixed database role can perform all configuration and maintenance activities on the database, and can also drop the database in SQL Server.
D: CREATE USER adds a user to the current database.
Note: Logins are created at the server level, while users are created at the database level. In other words, a login allows you to connect to the SQL Server service (also called an instance), and permissions inside the database are granted to the database users, not the logins. The logins will be assigned to server roles (for example, serveradmin) and the database users will be assigned to roles within that database (eg. db_datareader, db_bckupoperator).
References:
https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-role-transact-sql https://docs.microsoft.com/en-us/sql/t-sql/statements/create-user-transact-sql

NEW QUESTION 16
- (Exam Topic 3)
Your company has on-premises Microsoft SQL Server instance.
The data engineering team plans to implement a process that copies data from the SQL Server instance to Azure Blob storage. The process must orchestrate and manage the data lifecycle.
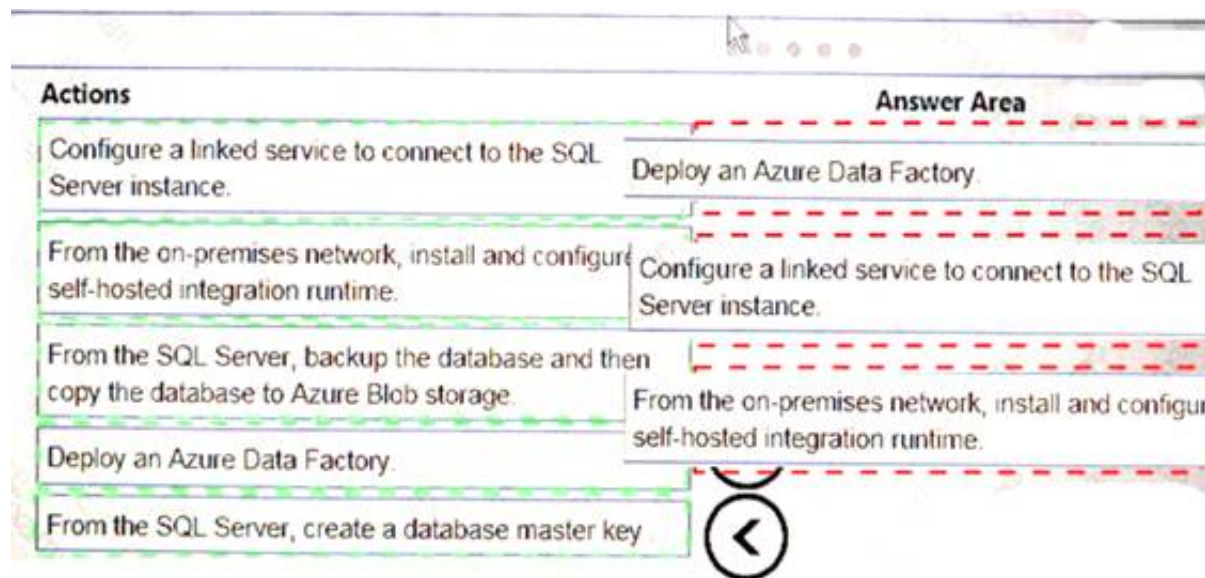You need to configure Azure Data Factory to connect to the SQL Server instance.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Actions	Answer Area
Configure a linked service to connect to the SQL Server instance.	Deploy an Azure Data Factory.
From the on-premises network, install and configure self-hosted integration runtime.	Configure a linked service to connect to the SQL Server instance.
From the SQL Server, backup the database and then copy the database to Azure Blob storage.	From the on-premises network, install and configure self-hosted integration runtime.
Deploy an Azure Data Factory.	
From the SQL Server, create a database master key	

NEW QUESTION 18
- (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

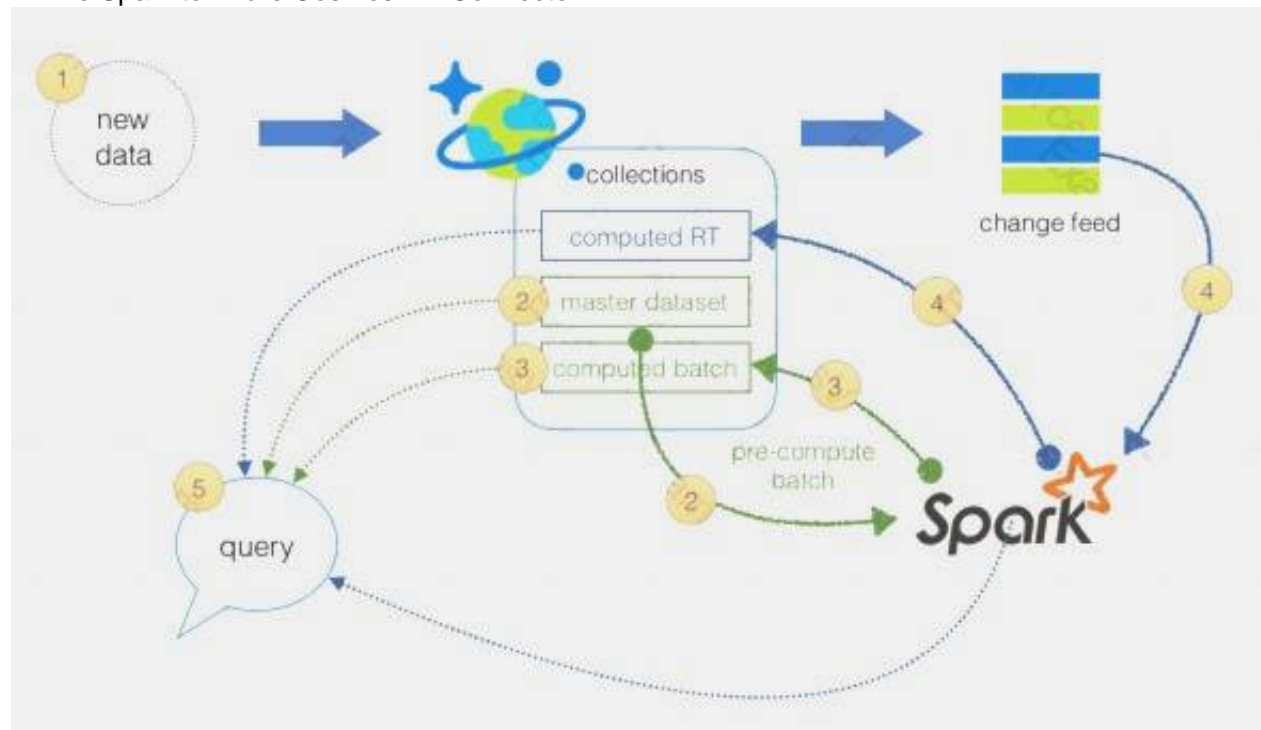**Answer:** D

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-
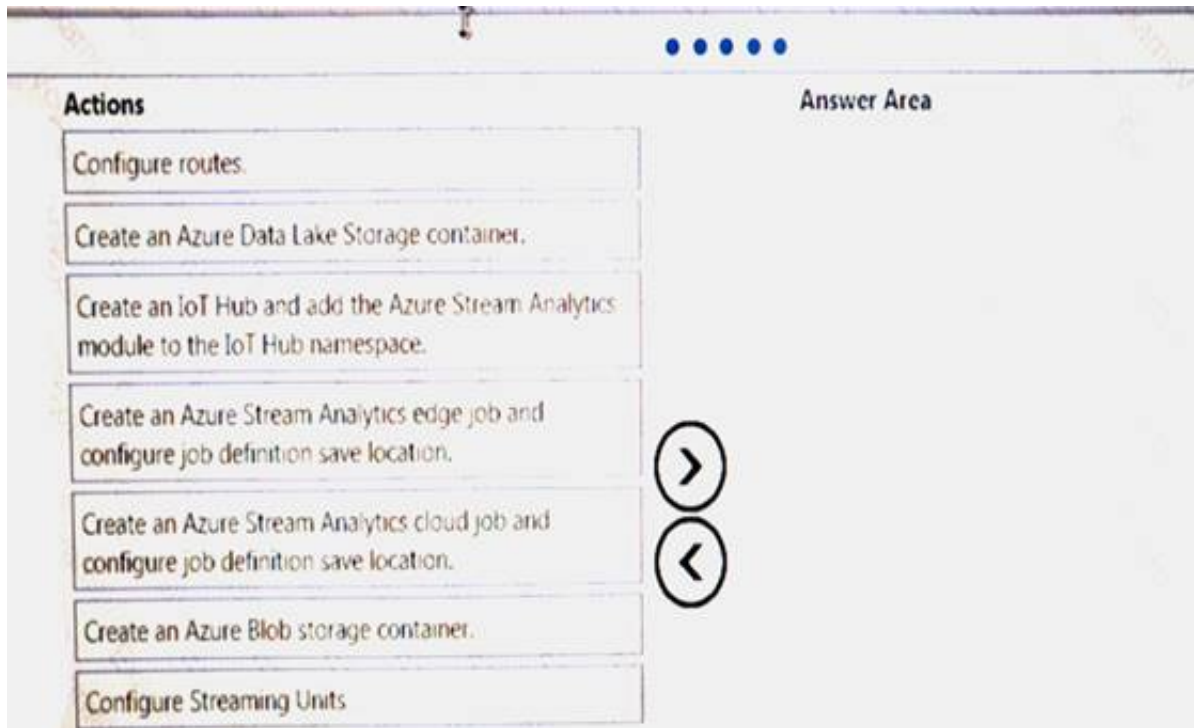
NEW QUESTION 22
- (Exam Topic 3)
You develop data engineering solutions for a company.
You need to deploy a Microsoft Azure Stream Analytics job for an IoT solution. The solution must:
• Minimize latency.
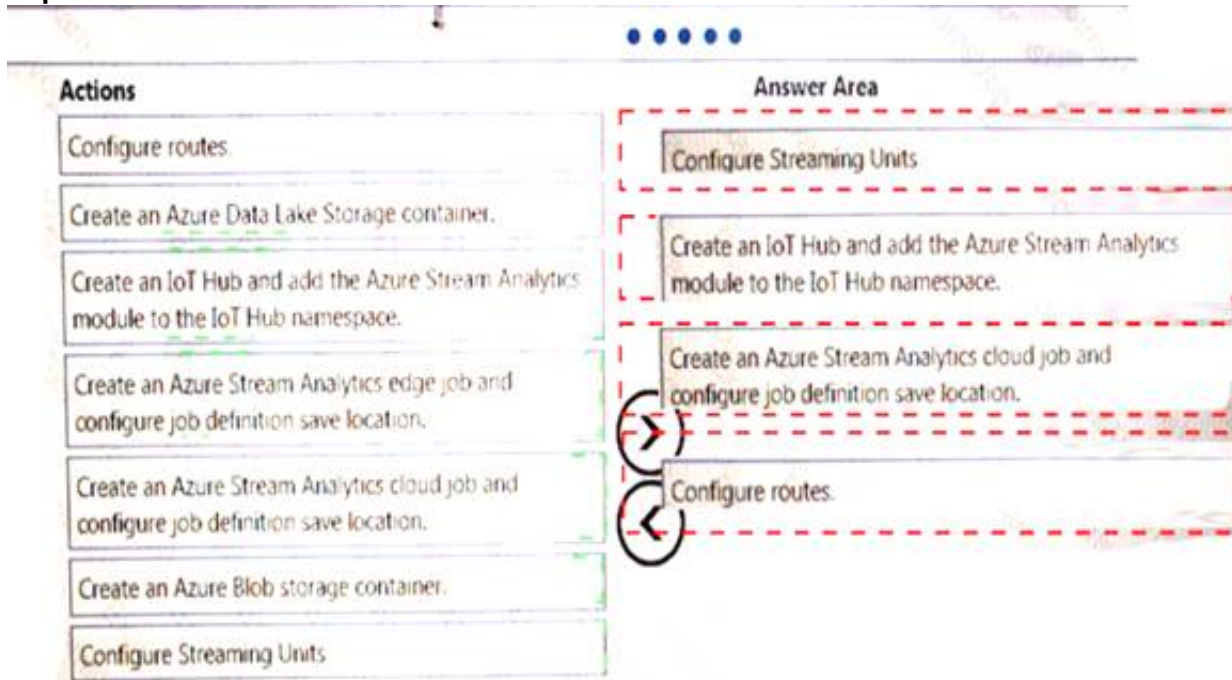• Minimize bandwidth usage between the job and IoT device.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 23
- (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.
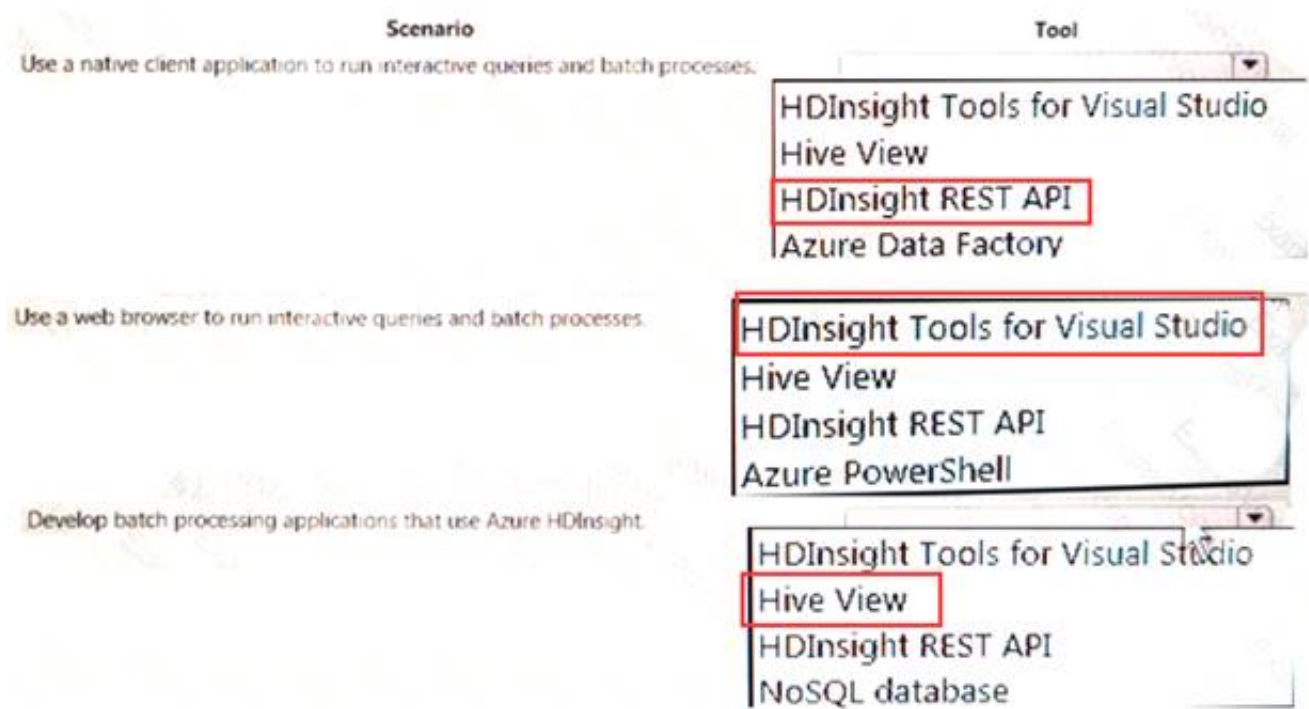Which Azure services should you use? To answer, select the appropriate configuration tit the answer area. NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 26
- (Exam Topic 3)
A company has a real-lime data analysis solution that is hosted on Microsoft Azure the solution uses Azure Event Hub to ingest data and an Azure Stream Analytics cloud job to analyze the data. The cloud job is configured to use 120 Streaming Units (SU).
You need to optimize performance for the Azure Stream Analytics job.
Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one port.

A. Implement event ordering
B. Scale the SU count for the job up
C. Implement Azure Stream Analytics user-defined functions (UDF)
D. Scale the SU count for the job down
E. Implement query parallelization by partitioning the data output
F. Implement query parallelization by partitioning the data input

Answer: BF

Explanation:
Scale out the query by allowing the system to process each input partition separately.
F: A Stream Analytics job definition includes inputs, a query, and output. Inputs are where the job reads the data stream from.
References:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-parallelization

NEW QUESTION 29
- (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 Studi
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 31
- (Exam Topic 3)
You develop data engineering solutions for a company.
A project requires analysis of real-time Twitter feeds. Posts that contain specific keywords must be stored and processed on Microsoft Azure and then displayed by using Microsoft Power BI. You need to implement the solution.
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.

Actions	Answer Area
Create an HDInsight cluster with the Hadoop cluster type.	
Create a Jupyter Notebook.	
Run a job that uses the Spark Streaming API to ingest data from Twitter.	
Create a Runbook.	
Create an HDInsight cluster with the Spark cluster type.	
Create an table.	
Load the hvac table into Power BI Desktop.	

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Create an HDInisght cluster with the Spark cluster type Step 2: Create a Jyputer Notebook
Step 3: Create a table
The Jupyter Notebook that you created in the previous step includes code to create an hvac table. Step 4: Run a job that uses the Spark Streaming API to ingest data from Twitter
Step 5: Load the hvac table into Power BI Desktop
You use Power BI to create visualizations, reports, and dashboards from the Spark cluster data. References:
https://acadgild.com/blog/streaming-twitter-data-using-spark
https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-use-with-data-lake-store

NEW QUESTION 36
- (Exam Topic 3)
An application will use Microsoft Azure Cosmos DB as its data solution. The application will use the Cassandra API to support a column-based database type that uses containers to store items.
You need to provision Azure Cosmos DB. Which container name and item name should you use? Each correct answer presents part of the solutions.
NOTE: Each correct answer selection is worth one point.

A. table
B. collection
C. graph
D. entities
E. rows

Answer: AE

Explanation:
Depending on the choice of the API, an Azure Cosmos item can represent either a document in a collection, a row in a table or a node/edge in a graph. The following table shows the mapping between API-specific entities to an Azure Cosmos item:

Cosmos entity	SQL API	Cassandra API	Azure Cosmos DB's API for MongoDB	Gremlin API	Table API
Azure Cosmos item	Document	Row	Document	Node or Edge	Item

An Azure Cosmos container is specialized into API-specific entities as follows:

Azure Cosmos entity	SQL API	Cassandra API	Azure Cosmos DB's API for MongoDB	Gremlin API	Table API
Azure Cosmos container	Collection	Table	Collection	Graph	Table

References:
https://docs.microsoft.com/en-us/azure/cosmos-db/databases-containers-items

NEW QUESTION 38
- (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 40
- (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 data engineering solutions for a company.
A project requires the deployment of resources to Microsoft Azure for batch data processing on Azure
HDInsight. Batch processing will run daily and must: Scale to minimize costs
Be monitored for cluster performance
You need to recommend a tool that will monitor clusters and provide information to suggest how to scale. Solution: Monitor cluster load using the Ambari Web UI.
Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Ambari Web UI does not provide information to suggest how to scale.
Instead monitor clusters by using Azure Log Analytics and HDInsight cluster management solutions. References:
https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hadoop-oms-log-analytics-tutorial https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hadoop-manage-ambari

NEW QUESTION 45
- (Exam Topic 3)
You develop data engineering solutions for a company. An application creates a database on Microsoft Azure. You have the following code:
Which database and authorization types are used? To answer, select the appropriate option in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Component	Technology
Azure database type	Azure Cosmos DB Azure SQL Database files Blob ∨
Key type	resource token Master key certificate ∨

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Azure Cosmos DB
The DocumentClient.CreateDatabaseAsync(Database, RequestOptions) method creates a database resource as an asychronous operation in the Azure Cosmos DB service.
Box 2: Master Key
Azure Cosmos DB uses two types of keys to authenticate users and provide access to its data and resources: Master Key, Resource Tokens
Master keys provide access to the all the administrative resources for the database account. Master keys: ⊙ Provide access to accounts, databases, users, and permissions.
⊙ Cannot be used to provide granular access to containers and documents.
⊙ Are created during the creation of an account.
⊙ Can be regenerated at any time.

NEW QUESTION 49
- (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 51
- (Exam Topic 3)
A company is designing a hybrid solution to synchronize data and on-premises Microsoft SQL Server database to Azure SQL Database.
You must perform an assessment of databases to determine whether data will move without compatibility issues.
You need to perform the assessment. Which tool should you use?

A. Azure SQL Data Sync
B. SQL Vulnerability Assessment (VA)
C. SQL Server Migration Assistant (SSMA)
D. Microsoft Assessment and Planning Toolkit
E. Data Migration Assistant (DMA)

Answer: E

Explanation:
The Data Migration Assistant (DMA) helps you upgrade to a modern data platform by detecting compatibility issues that can impact database functionality in your new version of SQL Server or Azure SQL Database. DMA recommends performance and reliability improvements for your target environment and allows you to move your schema, data, and uncontained objects from your source server to your target server.
References:
https://docs.microsoft.com/en-us/sql/dma/dma-overview

NEW QUESTION 53
- (Exam Topic 3)
Your company uses Microsoft Azure SQL Database configure with Elastic pool. You use Elastic Database jobs to run queries across all databases in the pod.
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 services 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 56
- (Exam Topic 3)
You plan to create a new single database instance of Microsoft Azure SQL Database.
The database must only allow communication from the data engineer's workstation. You must connect directly to the instance by using Microsoft SQL Server Management Studio.
You need to create and configure the Database. Which three Azure PowerShell cmdlets should you use to develop the solution? To answer, move the appropriate cmdlets from the list of cmdlets to the answer area and arrange them in the correct order.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: New-AzureSqlServer Create a server.
Step 2: New-AzureRmSqlServerFirewallRule
New-AzureRmSqlServerFirewallRule creates a firewall rule for a SQL Database server. Can be used to create a server firewall rule that allows access from the specified IP range. Step 3: New-AzureRmSqlDatabase
Example: Create a database on a specified server
PS C:\>New-AzureRmSqlDatabase -ResourceGroupName "ResourceGroup01" -ServerName "Server01"
-DatabaseName "Database01
References:
https://docs.microsoft.com/en-us/azure/sql-database/scripts/sql-database-create-and-configure-database-powersh

NEW QUESTION 61

- (Exam Topic 3)
A company runs Microsoft Dynamics CRM with Microsoft SQL Server on-premises. SQL Server Integration Services (SSIS) packages extract data from Dynamics CRM APIs, and load the data into a SQL Server data warehouse.
The datacenter is running out of capacity. Because of the network configuration, you must extract on premises data to the cloud over https. You cannot open any additional ports. The solution must implement the least amount of effort.
You need to create the pipeline system.
Which component should you use? To answer, select the appropriate technology in the dialog box in the answer area.
NOTE: Each correct selection is worth one point.

Action	Technology
Extract SQL data on-premises	Self-hosted integration runtime / Azure-SSIS integration runtime / Azure integration runtime / Source ∨
Load SQL data warehouse	Self-hosted integration runtime / Azure-SSIS integration runtime / Azure integration runtime / Sink ∨

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Source
For Copy activity, it requires source and sink linked services to define the direction of data flow. Copying between a cloud data source and a data source in private network: if either source or sink linked
service points to a self-hosted IR, the copy activity is executed on that self-hosted Integration Runtime.
Box 2: Self-hosted integration runtime
A self-hosted integration runtime can run copy activities between a cloud data store and a data store in a private network, and it can dispatch transform activities against compute resources in an on-premises network or an Azure virtual network. The installation of a self-hosted integration runtime needs on an on-premises machine or a virtual machine (VM) inside a private network.
References:
https://docs.microsoft.com/en-us/azure/data-factory/create-self-hosted-integration-runtime

NEW QUESTION 63
- (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 66
- (Exam Topic 3)
You develop data engineering solutions for a company.
A project requires the deployment of data to Azure Data Lake Storage.
You need to implement role-based access control (RBAC) so that project members can manage the Azure Data Lake Storage resources.

Which three actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Assign Azure AD security groups to Azure Data Lake Storage.
B. Configure end-user authentication for the Azure Data Lake Storage account.
C. Configure service-to-service authentication for the Azure Data Lake Storage account.
D. Create security groups in Azure Active Directory (Azure AD) and add project members.
E. Configure access control lists (ACL) for the Azure Data Lake Storage account.

Answer: ADE

NEW QUESTION 67
- (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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 69
- (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 data engineering solutions for a company.
A project requires the deployment of resources to Microsoft Azure for batch data processing on Azure
HDInsight. Batch processing will run daily and must: Scale to minimize costs
Be monitored for cluster performance
You need to recommend a tool that will monitor clusters and provide information to suggest how to scale. Solution: Download Azure HDInsight cluster logs by using Azure PowerShell.
Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Reference:
Instead monitor clusters by using Azure Log Analytics and HDInsight cluster management solutions. References:
https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hadoop-oms-log-analytics-tutorial

NEW QUESTION 70
- (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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 74
- (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 Data Lake Gen 2 storage account.
2. Create an external tile format and external table using the external data source.
3. Load the data using the CREATE TABLE AS SELECT statement.
Does the solution meet the goal?

A. Yes
B. No

Answer: A

NEW QUESTION 75
- (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 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.

A company uses Azure Data Lake Gen 1 Storage to store big data related to consumer behavior. You need to implement logging.
Solution: Use information stored m Azure Active Directory reports.
Does the solution meet the goal?

A. Yes
B. No

Answer: B

NEW QUESTION 77
- (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 82
- (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 toad 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 an external file format and external table using the external data source
3. Load the data using the INSERT…SELECT statement
Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
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 83
......

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your DP-200 Exam with Our Prep Materials Via below:

https://www.certleader.com/DP-200-dumps.html