DP-200 Dumps

Implementing an Azure Data Solution

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

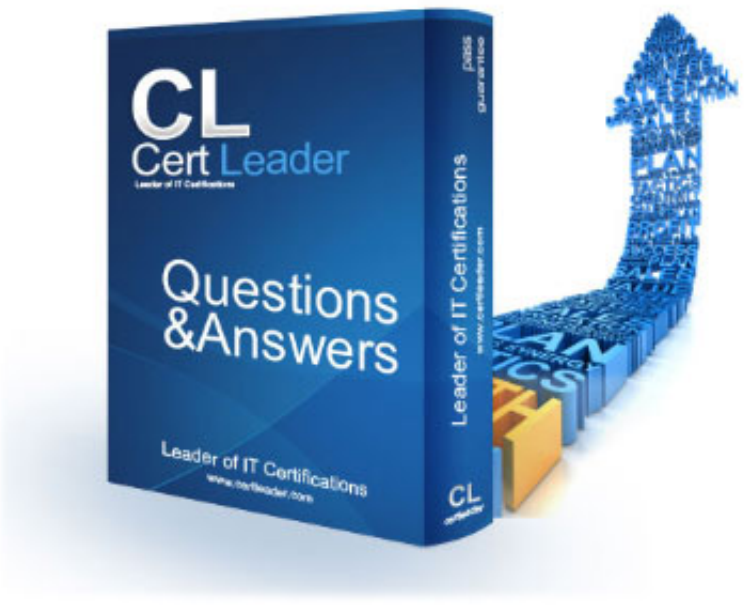

NEW QUESTION 1
- (Exam Topic 1)
You need to provision the polling data storage account.
How should you configure the storage account? To answer, drag the appropriate Configuration Value to the correct Setting. Each Configuration Value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

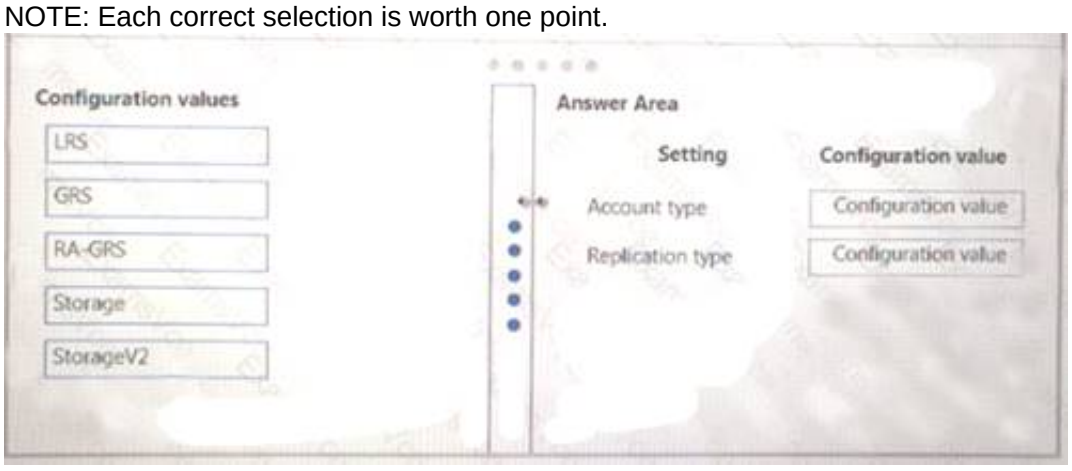

A. Mastered
B. Not Mastered

Answer: A

Explanation:

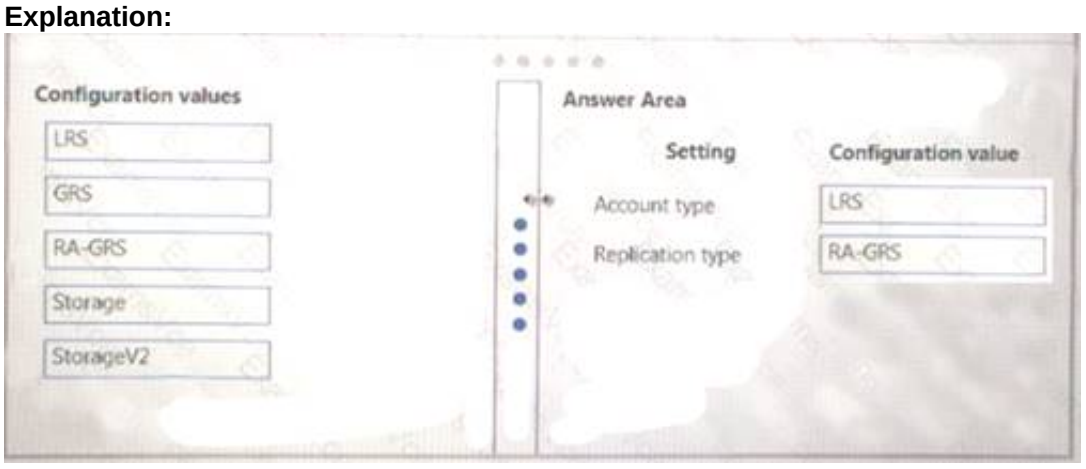

NEW QUESTION 2
- (Exam Topic 1)
You need to ensure that phone-based polling data can be analyzed in the PollingData database.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer are and arrange them in the correct order.

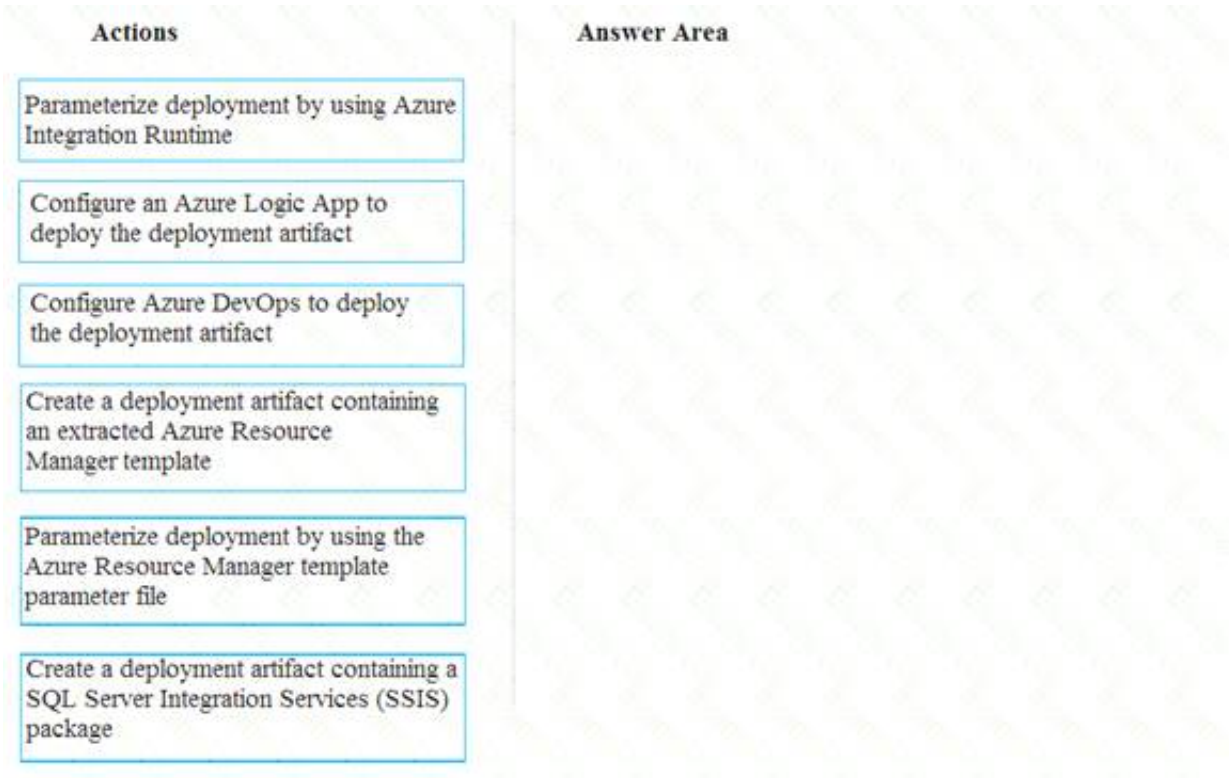

A. Mastered
B. Not Mastered

Answer: A

Explanation:

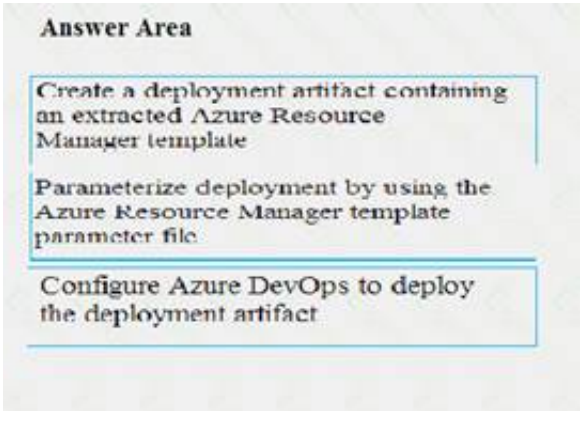

Scenario:
All deployments must be performed by using Azure DevOps. Deployments must use templates used in multiple environments
No credentials or secrets should be used during deployments

NEW QUESTION 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.

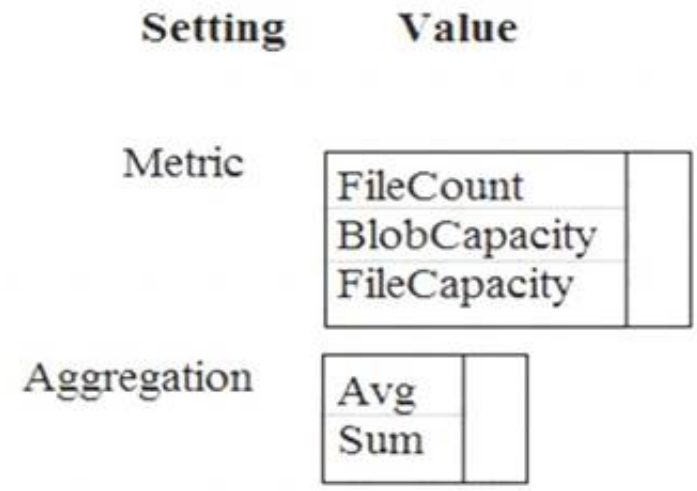

A. Mastered
B. Not Mastered

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 process and query ingested Tier 9 data.
Which two options should you use? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Azure Notification Hub
B. Transact-SQL statements
C. Azure Cache for Redis
D. Apache Kafka statements
E. Azure Event Grid
F. Azure Stream Analytics

Answer: EF

Explanation:
Event Hubs provides a Kafka endpoint that can be used by your existing Kafka based applications as an alternative to running your own Kafka cluster.
You can stream data into Kafka-enabled Event Hubs and process it with Azure Stream Analytics, in the following steps:
- Create a Kafka enabled Event Hubs namespace.
- Create a Kafka client that sends messages to the event hub.
- Create a Stream Analytics job that copies data from the event hub into an Azure blob storage. Scenario:

Internal Distribution and Sales	9	Yes, once ingested at branches	Data ingested from Contoso branches

Tier 9 reporting must be moved to Event Hubs, queried, and persisted in the same Azure region as the company's main office
References:
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-kafka-stream-analytics

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).

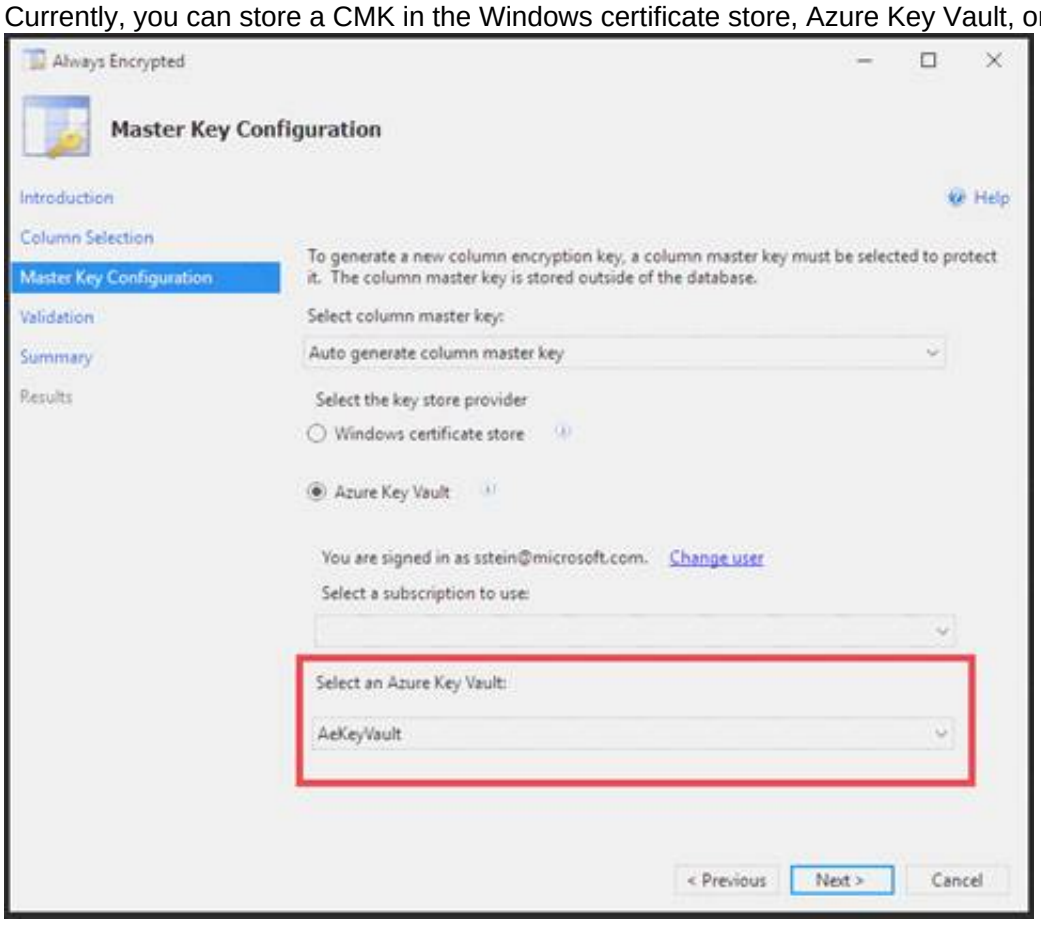

References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-always-encrypted-azure-key-vault

NEW QUESTION 6
- (Exam Topic 2)
You need to set up Azure Data Factory pipelines to meet data movement requirements. Which integration runtime should you use?

A. self-hosted integration runtime
B. Azure-SSIS Integration Runtime
C. .NET Common Language Runtime (CLR)
D. Azure integration runtime

Answer: A

Explanation:
The following table describes the capabilities and network support for each of the integration runtime types:

IR type	Public network	Private network
Azure	Data movement Activity dispatch	
Self-hosted	Data movement Activity dispatch	Data movement Activity dispatch
Azure-SSIS	SSIS package execution	SSIS package execution

Scenario: The solution must support migrating databases that support external and internal application to Azure SQL Database. The migrated databases will be supported by Azure Data Factory pipelines for the continued movement, migration and updating of data both in the cloud and from local core business systems and repositories.
References:
https://docs.microsoft.com/en-us/azure/data-factory/concepts-integration-runtime

NEW QUESTION 7
- (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 8
- (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.

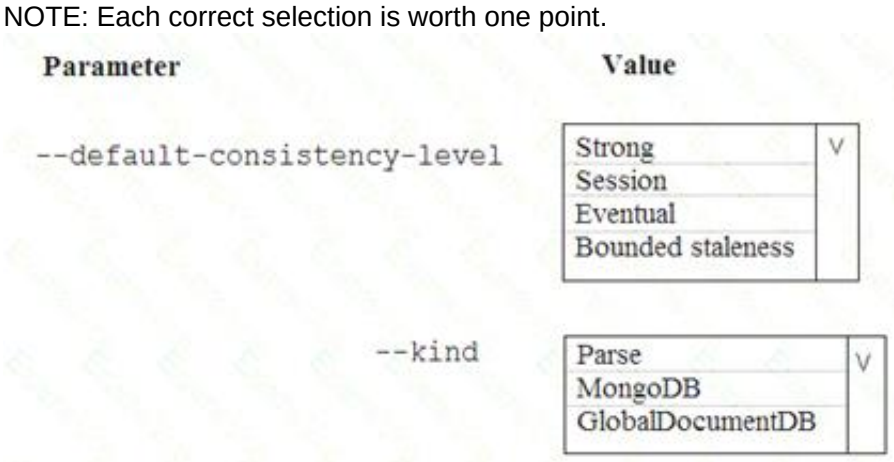

A. Mastered
B. Not Mastered

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 9
- (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 10

- (Exam Topic 3)
You develop data engineering solutions for a company. The company has on-premises Microsoft SQL Server databases at multiple locations.
The company must integrate data with Microsoft Power BI and Microsoft Azure Logic Apps. The solution must avoid single points of failure during connection and transfer to the cloud. The solution must also minimize latency.
You need to secure the transfer of data between on-premises databases and Microsoft Azure.
What should you do?

A. Install a standalone on-premises Azure data gateway at each location
B. Install an on-premises data gateway in personal mode at each location
C. Install an Azure on-premises data gateway at the primary location
D. Install an Azure on-premises data gateway as a cluster at each location

Answer: D

Explanation:
You can create high availability clusters of On-premises data gateway installations, to ensure your organization can access on-premises data resources used in Power BI reports and dashboards. Such clusters allow gateway administrators to group gateways to avoid single points of failure in accessing on-premises data resources. The Power BI service always uses the primary gateway in the cluster, unless it's not available. In that case, the service switches to the next gateway in the cluster, and so on.
References:
https://docs.microsoft.com/en-us/power-bi/service-gateway-high-availability-clusters

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 13
- (Exam Topic 3)
A company uses Microsoft Azure SQL Database to store sensitive company data. You encrypt the data and only allow access to specified users from specified locations.
You must monitor data usage, and data copied from the system to prevent data leakage.
You need to configure Azure SQL Database to email a specific user when data leakage occurs.
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.

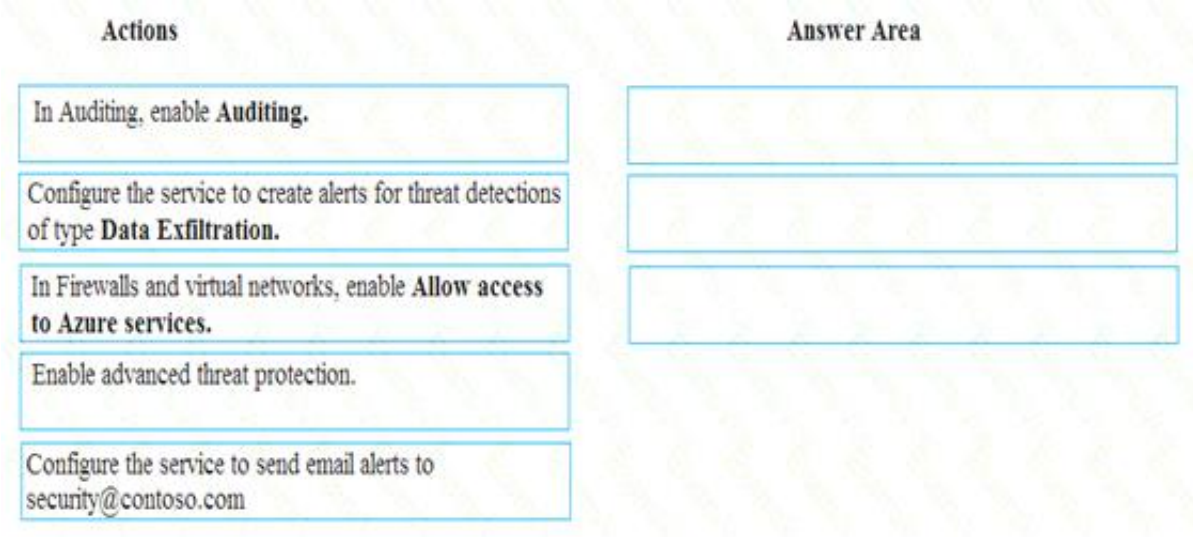

A. Mastered
B. Not Mastered

Answer: A

Explanation:

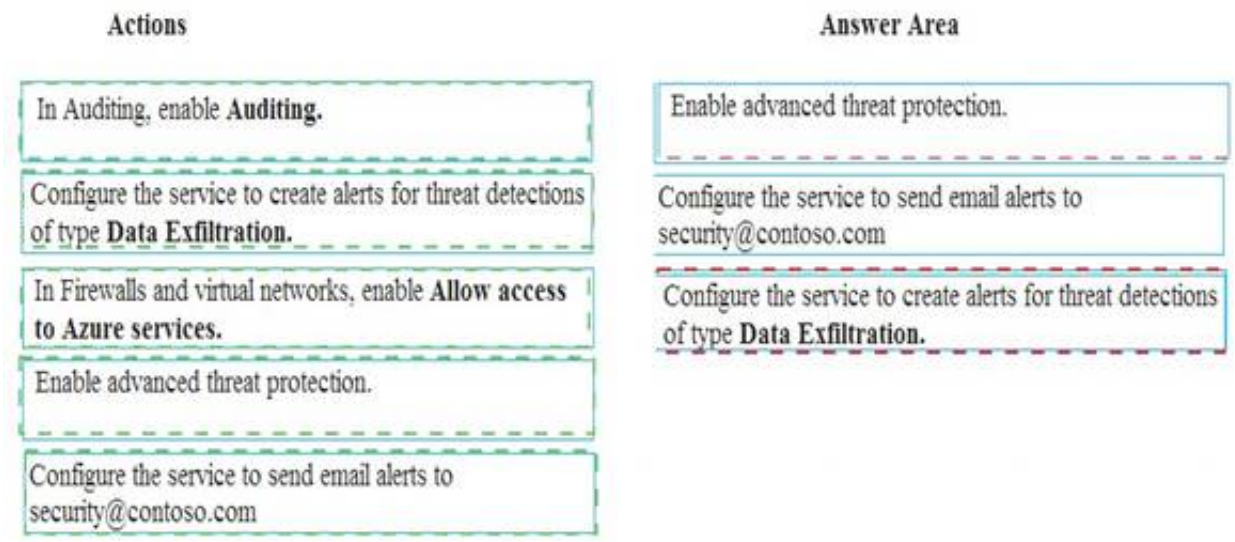

NEW QUESTION 16
- (Exam Topic 3)
You are developing the data platform for a global retail company. The company operates during normal working hours in each region. The analytical database is used once a week for building sales projections.
Each region maintains its own private virtual network.
Building the sales projections is very resource intensive are generates upwards of 20 terabytes (TB) of data. Microsoft Azure SQL Databases must be provisioned.
- Database provisioning must maximize performance and minimize cost
- The daily sales for each region must be stored in an Azure SQL Database instance
- Once a day, the data for all regions must be loaded in an analytical Azure SQL Database instance You need to provision Azure SQL database instances.

How should you provision the database instances? To answer, drag the appropriate Azure SQL products to the correct databases. Each Azure SQL product 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.

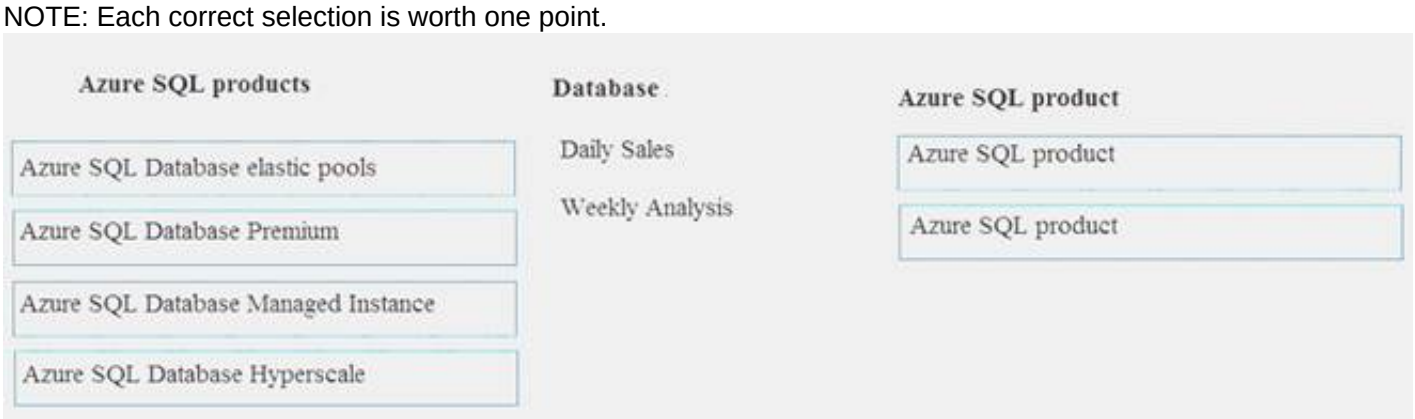

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Azure SQL Database elastic pools
SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single Azure
SQL Database server and share a set number of resources at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.
Box 2: Azure SQL Database Hyperscale
A Hyperscale database is an Azure SQL database in the Hyperscale service tier that is backed by the Hyperscale scale-out storage technology. A Hyperscale database supports up to 100 TB of data and provides high throughput and performance, as well as rapid scaling to adapt to the workload requirements. Scaling is transparent to the application – connectivity, query processing, and so on, work like any other SQL database.

NEW QUESTION 20
- (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.

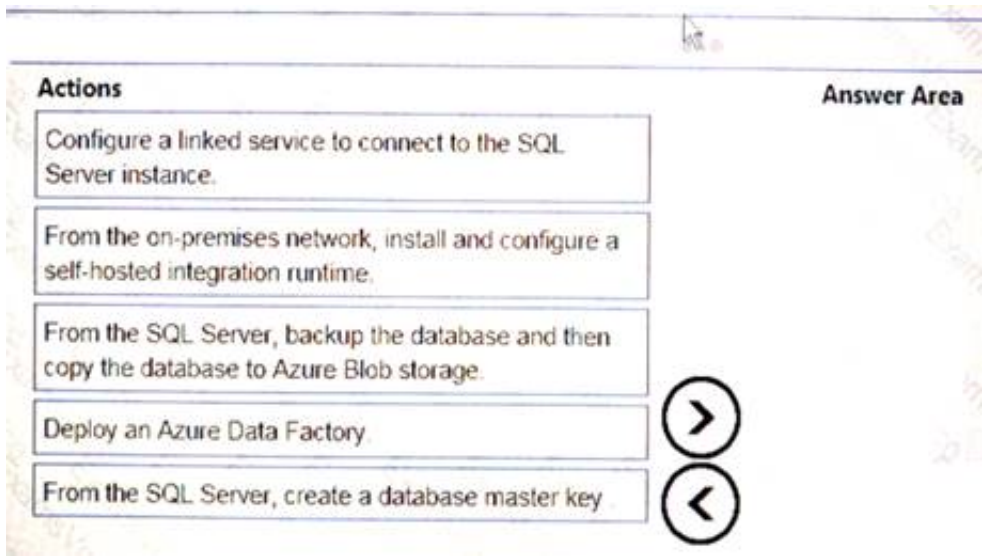

A. Mastered
B. Not Mastered

Answer: A

Explanation:

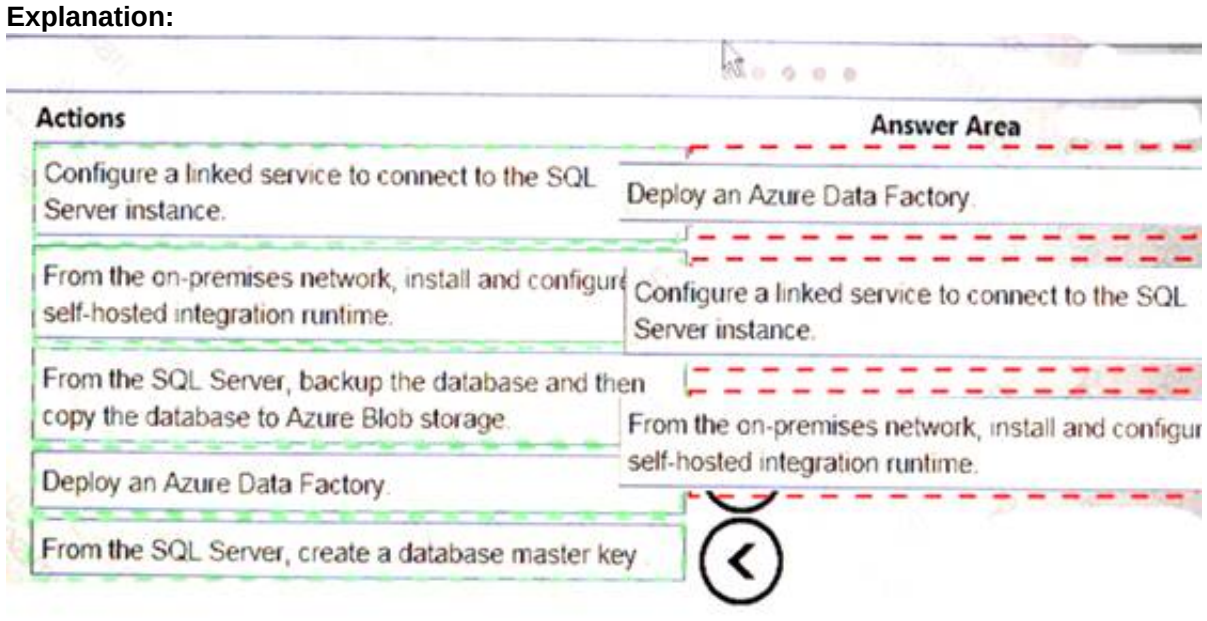

NEW QUESTION 24
- (Exam Topic 3)
Your company uses several Azure HDInsight clusters.
The data engineering team reports several errors with some application using these clusters. You need to recommend a solution to review the health of the clusters.
What should you include in you recommendation?

A. Azure Automation
B. Log Analytics
C. Application Insights

Answer: C

NEW QUESTION 29
- (Exam Topic 3)
A company runs Microsoft SQL Server in an on-premises virtual machine (VM).
You must migrate the database to Azure SQL Database. You synchronize users from Active Directory to Azure Active Directory (Azure AD).
You need to configure Azure SQL Database to use an Azure AD user as administrator. What should you configure?

A. For each Azure SQL Database, set the Access Control to administrator.
B. For the Azure SQL Database server, set the Active Directory to administrator.
C. For each Azure SQL Database, set the Active Directory administrator role.
D. For the Azure SQL Database server, set the Access Control to administrator.

Answer: A

NEW QUESTION 32
- (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

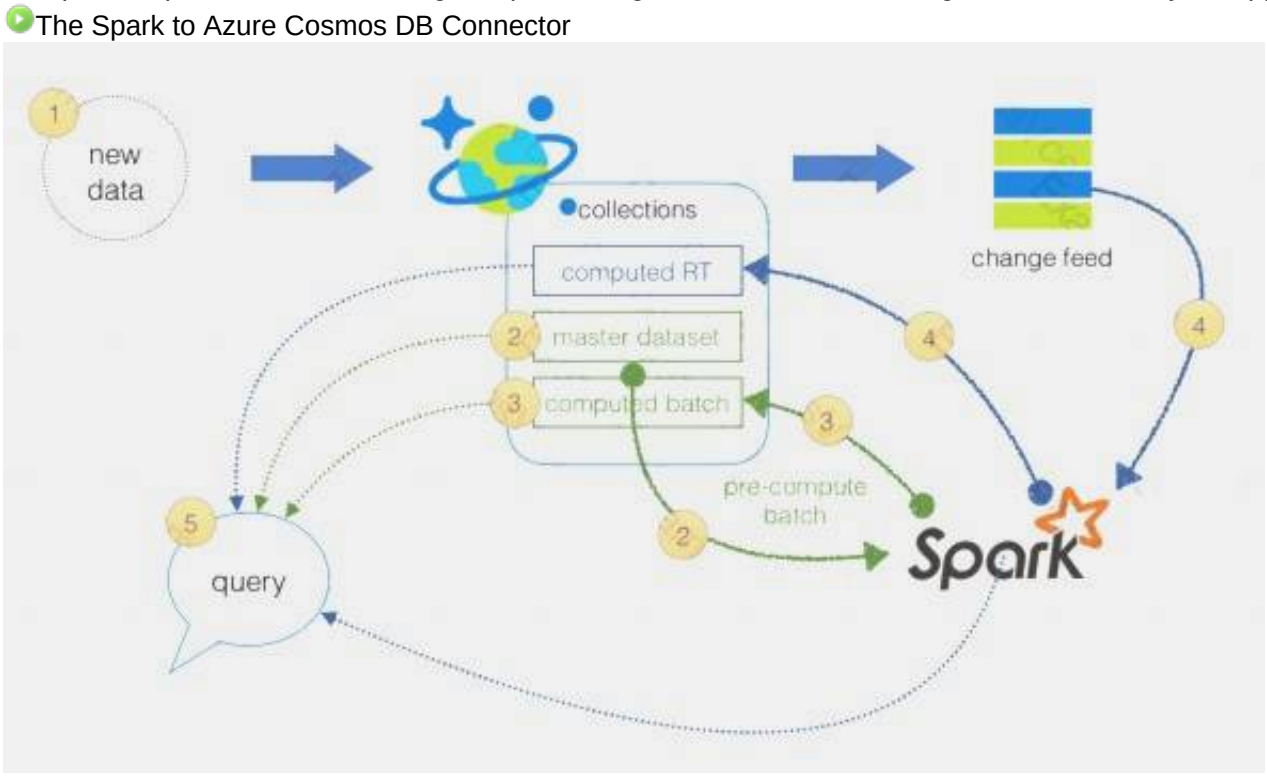

Note: Lambda architecture is a data-processing architecture designed to handle massive quantities of data by taking advantage of both batch processing and stream processing methods, and minimizing the latency involved in querying big data.
References:
https://sqlwithmanoj.com/2018/02/16/what-is-lambda-architecture-and-what-azure-offers-with-its-new-cosmos-

NEW QUESTION 37
- (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:

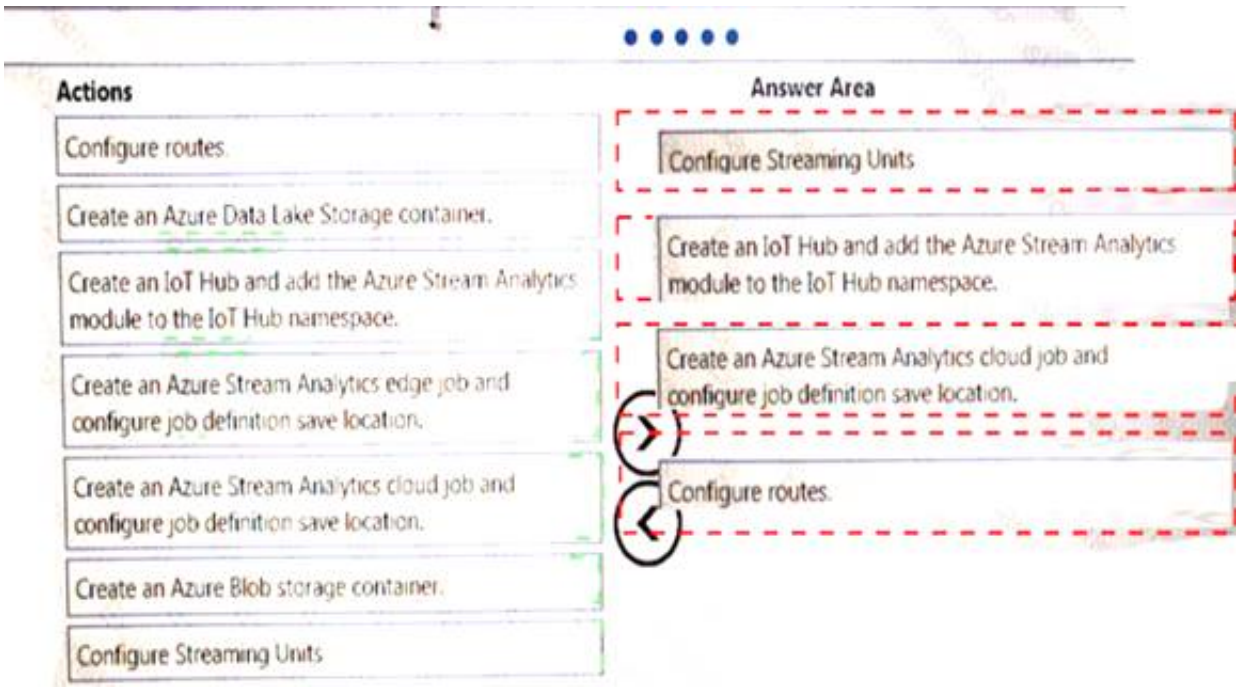

NEW QUESTION 38
- (Exam Topic 3)
A company has a real-lime data analysis solution that is hosted on Microsoft Azure the solution uses Azure Event Hub to ingest data and an Azure Stream Analytics cloud job to analyze the data. The cloud job is configured to use 120 Streaming Units (SU).
You need to optimize 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 40
- (Exam Topic 3)
A company is deploying a service-based data environment. You are developing a solution to process this data. The solution must meet the following requirements:
- Use an Azure HDInsight cluster for data ingestion from a relational database in a different cloud service
- Use an Azure Data Lake Storage account to store processed data
- Allow users to download processed data

You need to recommend technologies for the solution.
Which technologies should you use? To answer, select the appropriate options in the answer area.

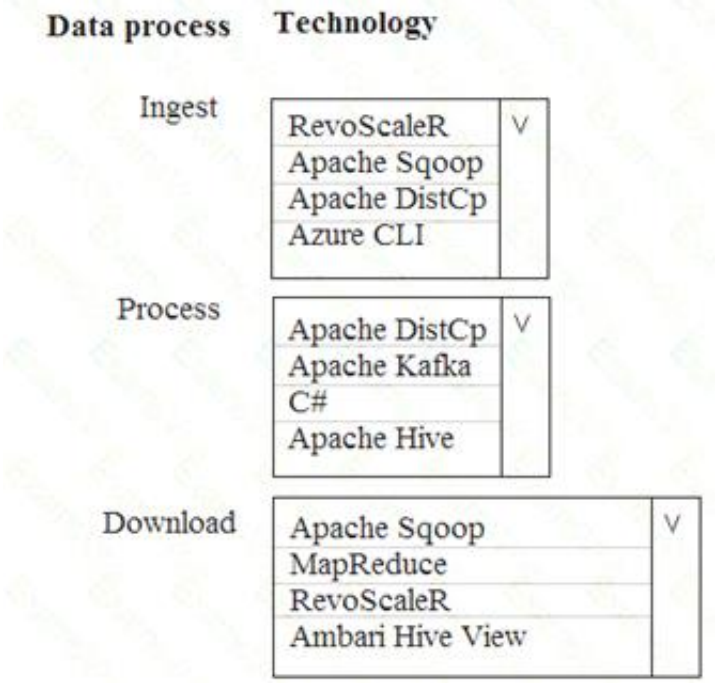

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Apache Sqoop is a tool designed for efficiently transferring bulk data between Apache Hadoop and structured datastores such as relational databases.

Azure HDInsight is a cloud distribution of the Hadoop components from the Hortonworks Data Platform (HDP).

NEW QUESTION 45
- (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 clusters by using Azure Log Analytics and HDInsight cluster management solutions. Does the solution meet the goal?

A. Yes
B. No

Answer: A

Explanation:
HDInsight provides cluster-specific management solutions that you can add for Azure Monitor logs. Management solutions add functionality to Azure Monitor logs, providing additional data and analysis tools. These solutions collect important performance metrics from your HDInsight clusters and provide the tools to search the metrics. These solutions also provide visualizations and dashboards for most cluster types
supported in HDInsight. By using the metrics that you collect with the solution, you can create custom monitoring rules and alerts.

NEW QUESTION 50
- (Exam Topic 3)
You need to develop a pipeline for processing data. The pipeline must meet the following requirements.
•Scale up and down resources for cost reduction.
•Use an in-memory data processing engine to speed up ETL and machine learning operations.
•Use streaming capabilities.
•Provide the ability to code in SQL, Python, Scala, and R.
•Integrate workspace collaboration with Git. What should you use?

A. HDInsight Spark Cluster
B. Azure Stream Analytics
C. HDInsight Hadoop Cluster
D. Azure SQL Data Warehouse

Answer: B

NEW QUESTION 53
- (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: Create an Azure Automation runbook to copy events. Does the solution meet the goal?

A. Yes
B. No

Answer: B

NEW QUESTION 56
- (Exam Topic 3)
A company manages several on-premises Microsoft SQL Server databases.
You need to migrate the databases to Microsoft Azure by using the backup process of Microsoft SQL Server. Which data technology should you use?

A. Azure SQL Database Managed Instance
B. Azure SQL Data Warehouse
C. Azure Cosmos DB
D. Azure SQL Database single database

Answer: D

NEW QUESTION 59
- (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.

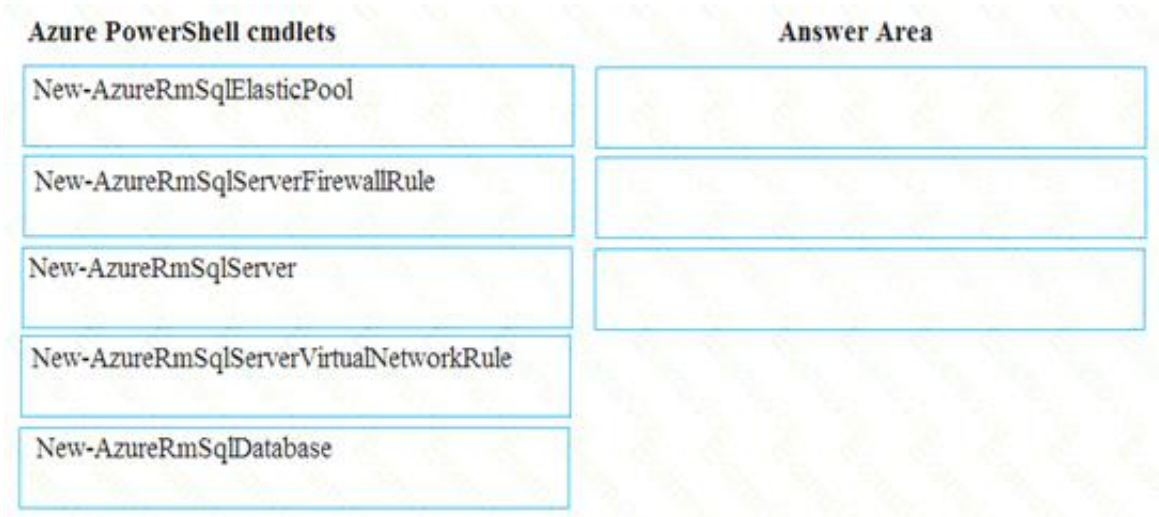

A. Mastered
B. Not Mastered

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 63
- (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.

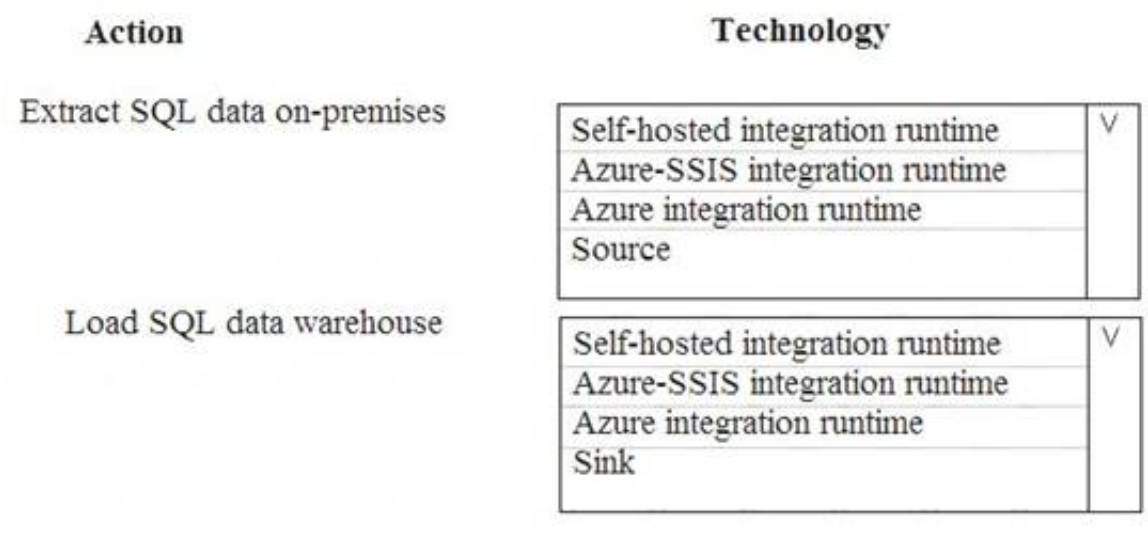

A. Mastered
B. Not Mastered

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 64
- (Exam Topic 3)
You manage the Microsoft Azure Databricks environment for a company. You must be able to access a private Azure Blob Storage account. Data must be available to all Azure Databricks workspaces. You need to provide the data access.

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.

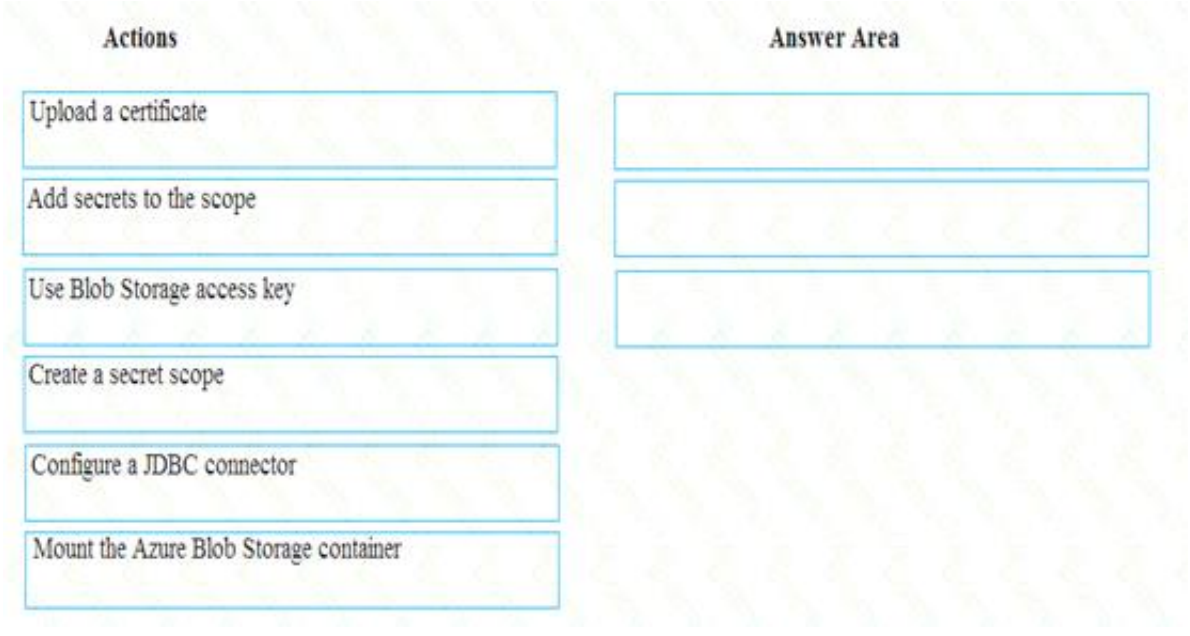

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Create a secret scope Step 2: Add secrets to the scope
Note: dbutils.secrets.get(scope = "<scope-name>", key = "<key-name>") gets the key that has been stored as a secret in a secret scope.
Step 3: Mount the Azure Blob Storage container
You can mount a Blob Storage container or a folder inside a container through Databricks File System - DBFS. The mount is a pointer to a Blob Storage container, so the data is never synced locally.
Note: To mount a Blob Storage container or a folder inside a container, use the following command:
Python dbutils.fs.mount(
source = "wasbs://<your-container-name>@<your-storage-account-name>.blob.core.windows.net", mount_point = "/mnt/<mount-name>",
extra_configs = {"<conf-key>":dbutils.secrets.get(scope = "<scope-name>", key = "<key-name>")}) where:
dbutils.secrets.get(scope = "<scope-name>", key = "<key-name>") gets the key that has been stored as a secret in a secret scope.
References:
https://docs.databricks.com/spark/latest/data-sources/azure/azure-storage.html

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 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 70
- (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.

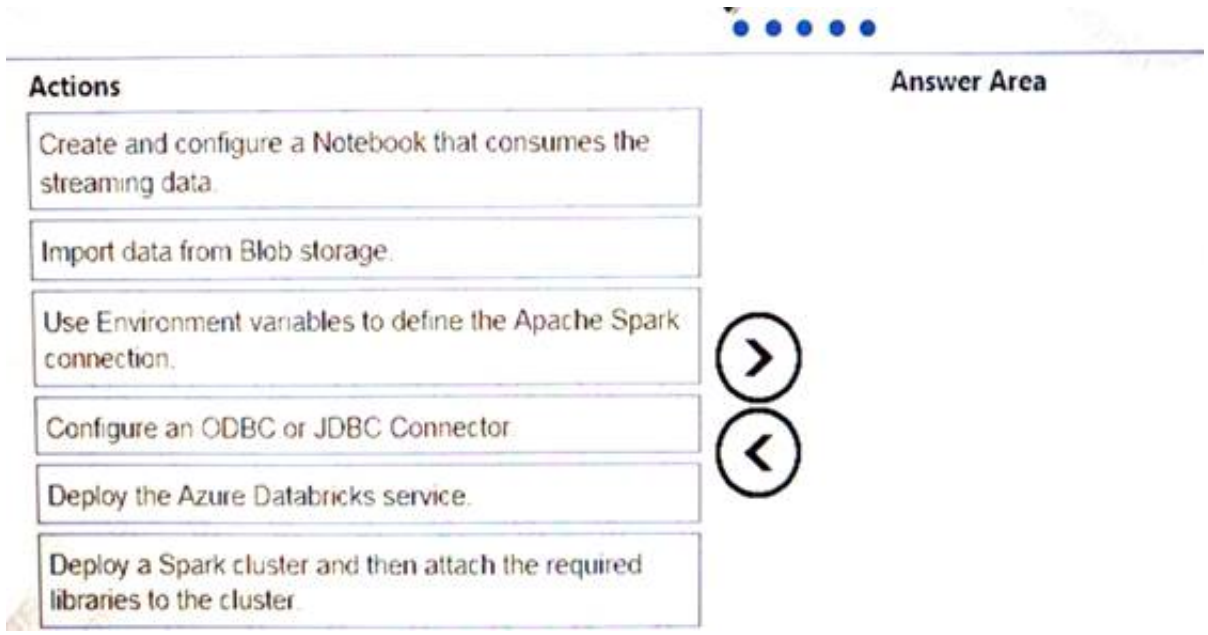

A. Mastered
B. Not Mastered

Answer: A

Explanation:

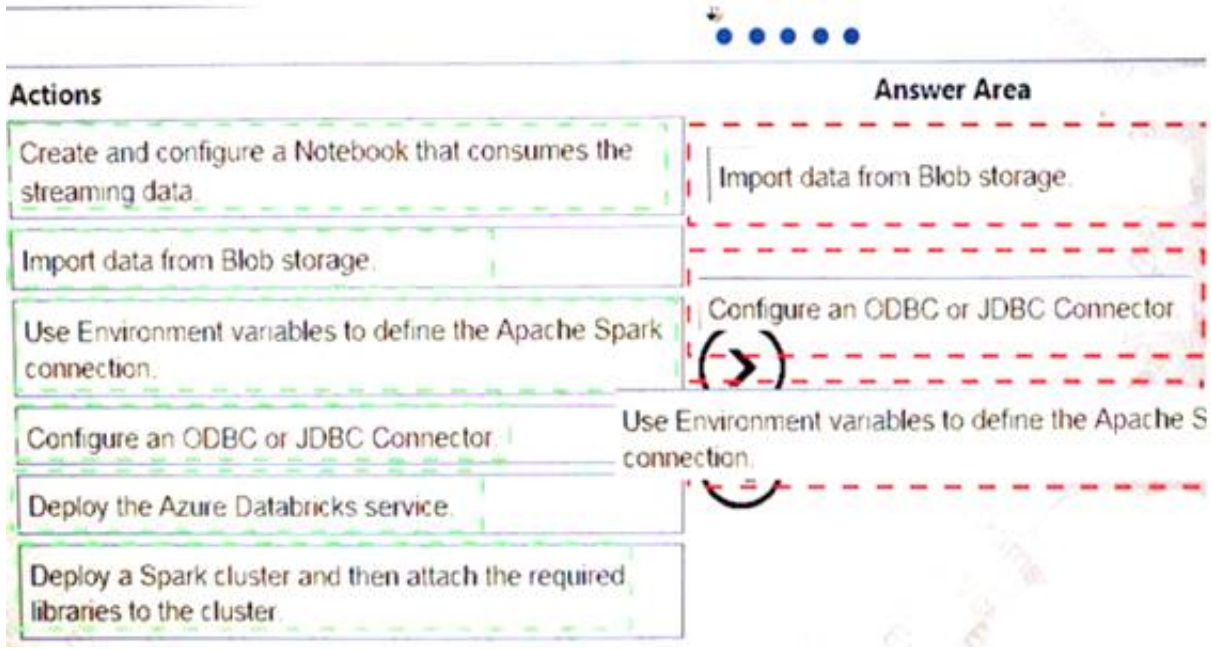

NEW QUESTION 72
- (Exam Topic 3)
You are creating a managed data warehouse solution on Microsoft Azure.
You must use PolyBase to retrieve data from Azure Blob storage that resides in parquet format and toad the data into a large table called FactSalesOrderDetails.
You need to configure Azure SQL Data 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:

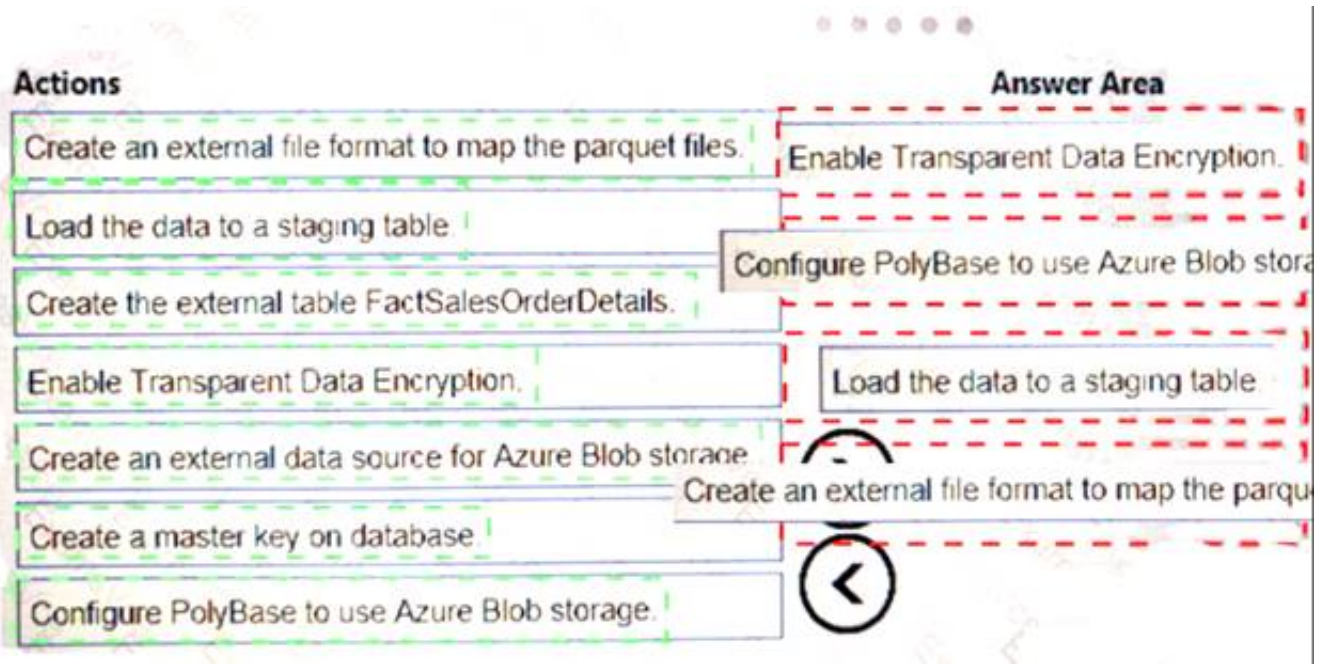

NEW QUESTION 76
- (Exam Topic 3)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.
You develop a data ingestion process that will import data to a Microsoft 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 80
- (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 85
- (Exam Topic 3)
A company has a Microsoft Azure HDInsight solution that uses different cluster types to process and analyze data. Operations are continuous. Reports indicate slowdowns during a specific lime window.
You need to determine a monitoring solution to track down the issue in the least amount of time. What should you use?

A. Azure Log Analytics log search query
B. Ambari REST API
C. Azure Monitor Metrics
D. HDInsight .NET SDK
E. Azure Log Analytics alert rule query

Answer: B

Explanation:
Ambari is the recommended tool for monitoring the health for any given HDInsight cluster.
Note: Azure HDInsight is a high-availability service that has redundant gateway nodes, head nodes, and ZooKeeper nodes to keep your HDInsight clusters running smoothly. While this ensures that a single failure will not affect the functionality of a cluster, you may still want to monitor cluster health so you are alerted when an issue does arise. Monitoring cluster health refers to monitoring whether all nodes in your cluster and the components that run on them are available and functioning correctly.
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.
References:
https://azure.microsoft.com/en-us/blog/monitoring-on-hdinsight-part-1-an-overview/

NEW QUESTION 89
- (Exam Topic 3)
You develop data engineering solutions for a company.
You need to ingest and visualize real-time Twitter data by using Microsoft Azure.
Which three technologies should you use? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Event Grid topic
B. Azure Stream Analytics Job that queries Twitter data from an Event Hub
C. Azure Stream Analytics Job that queries Twitter data from an Event Grid
D. Logic App that sends Twitter posts which have target keywords to Azure
E. Event Grid subscription
F. Event Hub instance

Answer: BDF

Explanation:
You can use Azure Logic apps to send tweets to an event hub and then use a Stream Analytics job to read from event hub and send them to PowerBI.
References:
https://community.powerbi.com/t5/Integrations-with-Files-and/Twitter-streaming-analytics-step-by-step/td-p/95

NEW QUESTION 90
.....

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