Microsoft

Exam Questions DP-200

Implementing an Azure Data Solution

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.

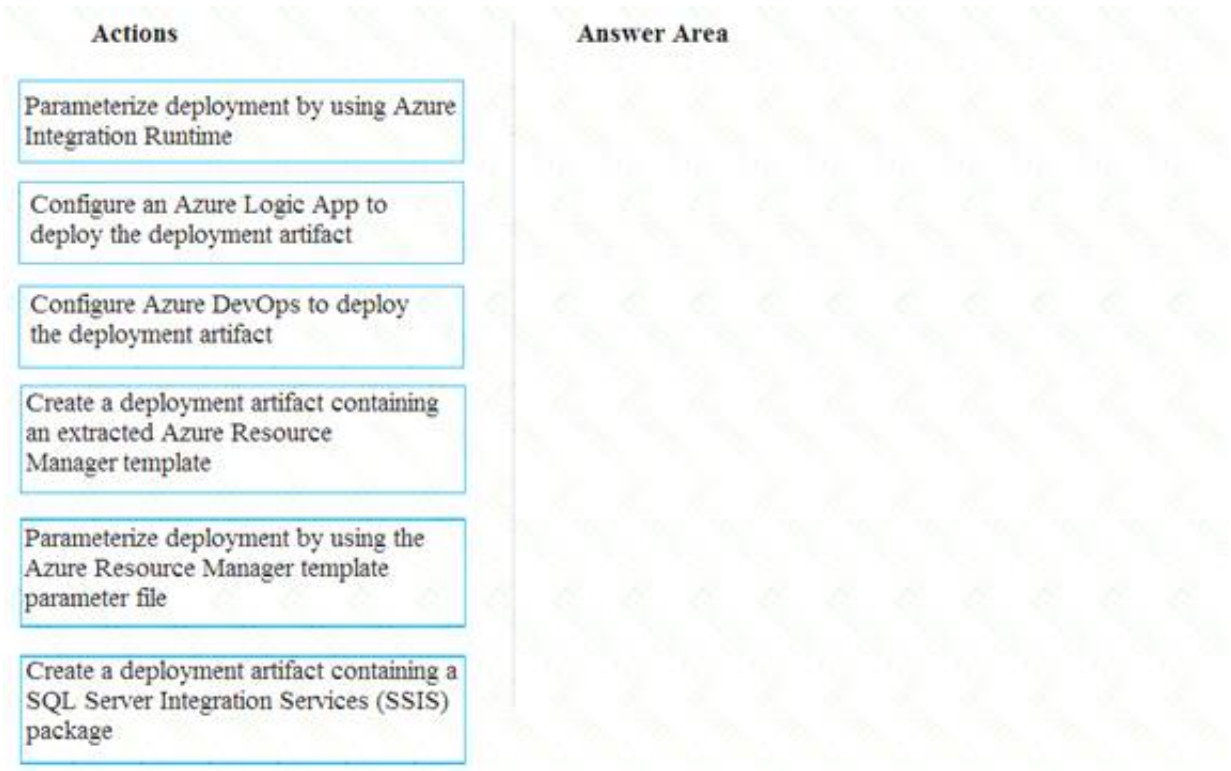

A. Mastered
B. Not Mastered

Answer: A

Explanation:

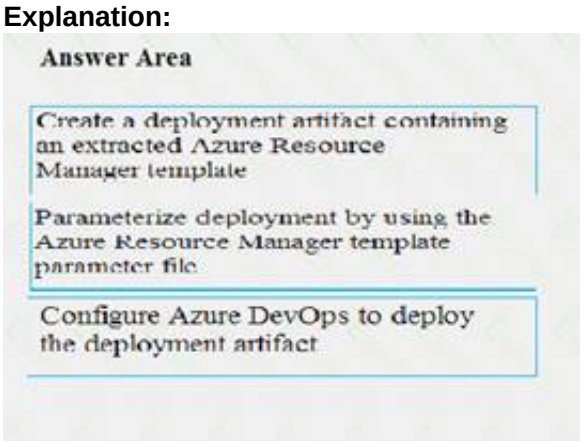

Scenario:
All deployments must be performed by using Azure DevOps. Deployments must use templates used in multiple environments
No credentials or secrets should be used during deployments

NEW QUESTION 2
- (Exam Topic 1)
You need to ensure that phone-based poling data can be analyzed in the PollingData database. How should you configure Azure Data Factory?

A. Use a tumbling schedule trigger
B. Use an event-based trigger
C. Use a schedule trigger
D. Use manual execution

Answer: C

Explanation:
When creating a schedule trigger, you specify a schedule (start date, recurrence, end date etc.) for the trigger, and associate with a Data Factory pipeline.
Scenario:
All data migration processes must use Azure Data Factory
All data migrations must run automatically during non-business hours References:
https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-schedule-trigger

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 implement diagnostic logging for Data Warehouse monitoring. Which log should you use?

A. RequestSteps
B. DmsWorkers

C. SqlRequests
D. ExecRequests

Answer: C

Explanation:
Scenario:
The Azure SQL Data Warehouse cache must be monitored when the database is being used.

Metric	Description
A	Low cache hit %, high cache usage %
B	Low cache hit %, low cache usage %
C	High cache hit %, high cache usage %

References:
https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-sql-r

NEW QUESTION 4
- (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).

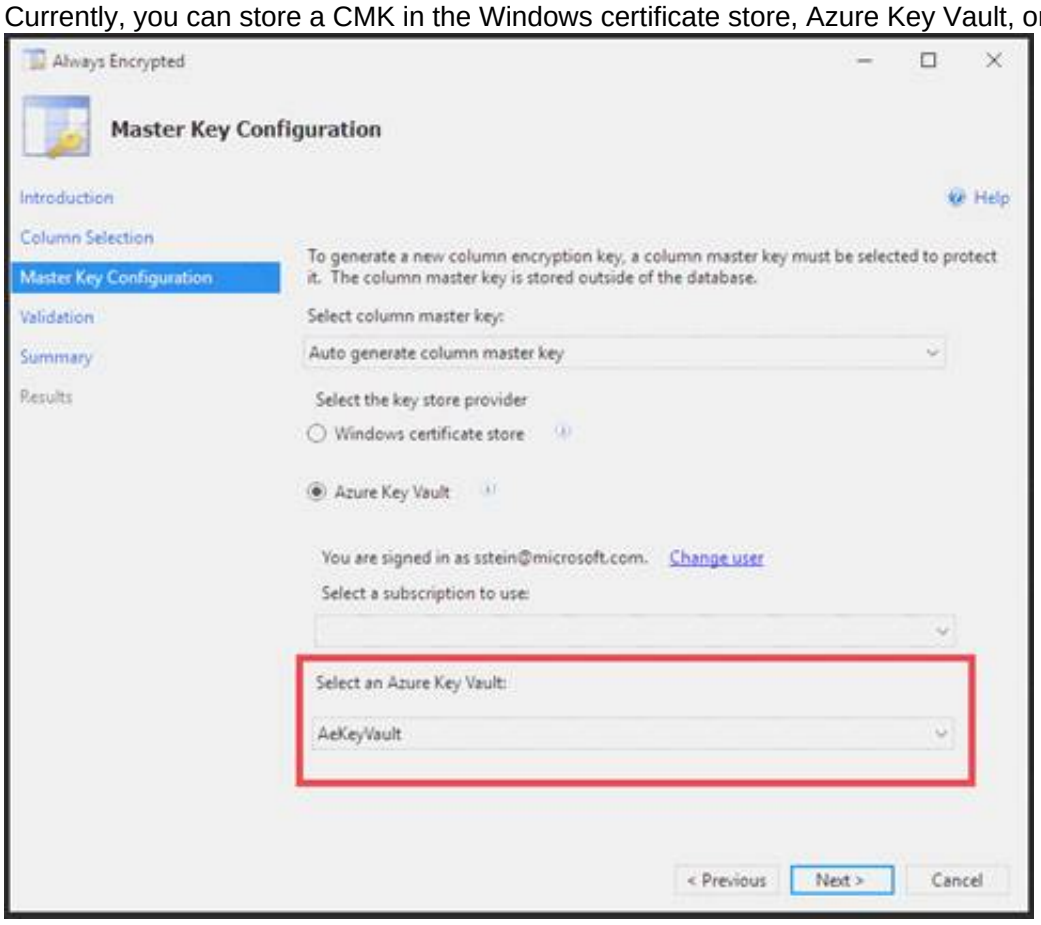

References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-always-encrypted-azure-key-vault

NEW QUESTION 5
- (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 6
- (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 7
- (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 8
- (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 9
- (Exam Topic 3)
You manage security for a database that supports a line of business application. Private and personal data stored in the database must be protected and encrypted. You need to configure the database to use Transparent Data Encryption (TDE).
Which five actions should you perform in sequence? To answer, select the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

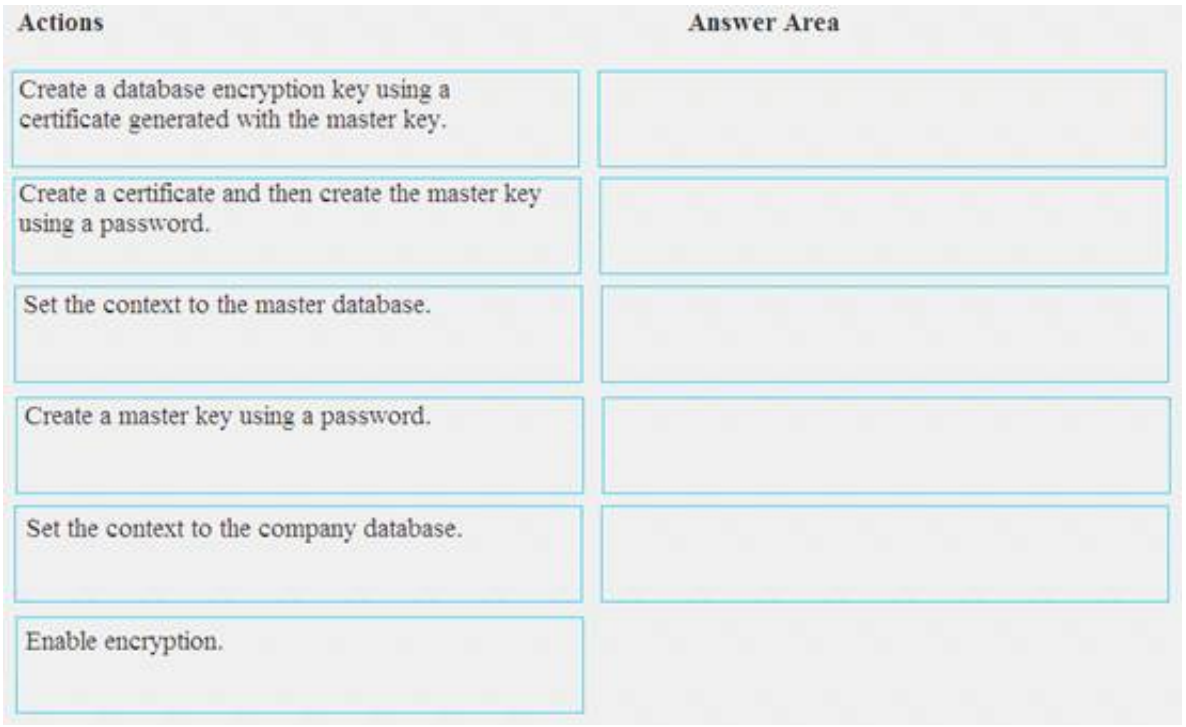

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Create a master key
Step 2: Create or obtain a certificate protected by the master key Step 3: Set the context to the company database
Step 4: Create a database encryption key and protect it by the certificate Step 5: Set the database to use encryption
Example code: USE master; GO
CREATE MASTER KEY ENCRYPTION BY PASSWORD = '<UseStrongPasswordHere>';
go
CREATE CERTIFICATE MyServerCert WITH SUBJECT = 'My DEK Certificate'; go
USE AdventureWorks2012; GO
CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = AES_128
ENCRYPTION BY SERVER CERTIFICATE MyServerCert; GO
ALTER DATABASE AdventureWorks2012 SET ENCRYPTION ON;
GO
References:
https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption

NEW QUESTION 10
- (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.

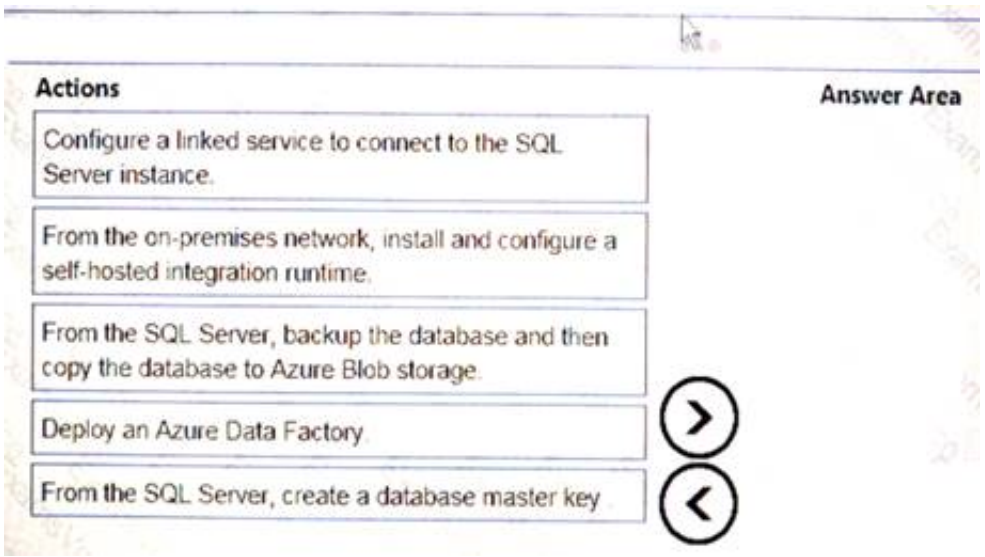

A. Mastered
B. Not Mastered

Answer: A

Explanation:

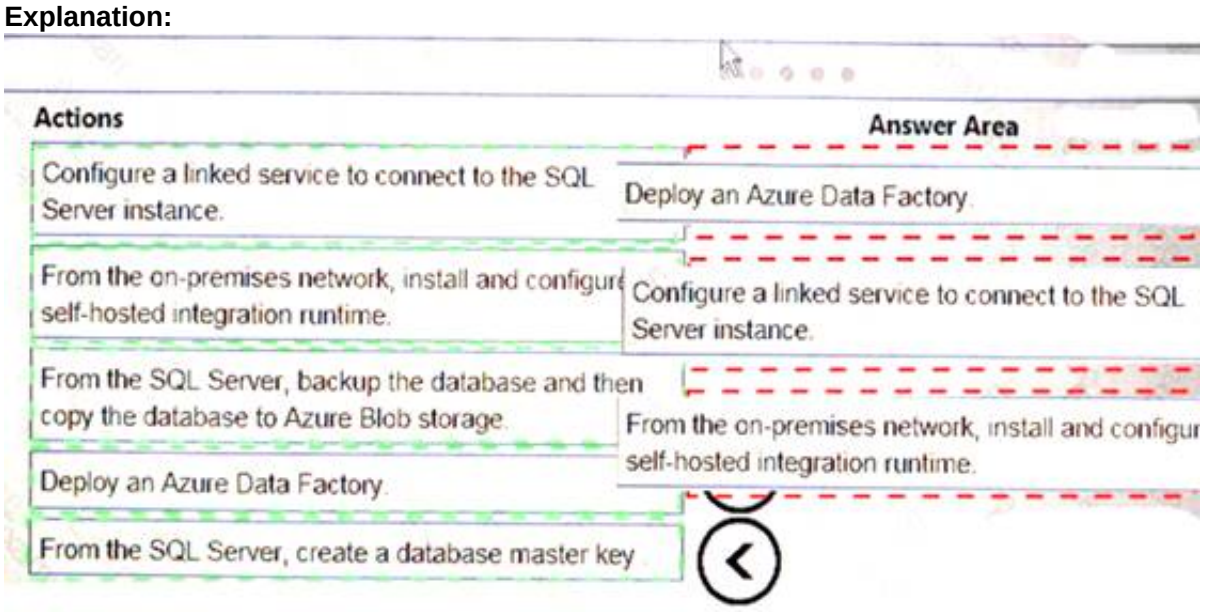

NEW QUESTION 10
- (Exam Topic 3)
Your company uses several Azure HDInsight clusters.
The data engineering team reports several errors with some application using these clusters. You need to recommend a solution to review the health of the clusters.
What should you include in you recommendation?

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

Answer: C

NEW QUESTION 14
- (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: Configure Azure Data Late Storage diagnostics to store logs and metrics in a storage account. Does the solution meet the goal?

A. Yes
B. No

Answer: A

NEW QUESTION 16
- (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 17
- (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:

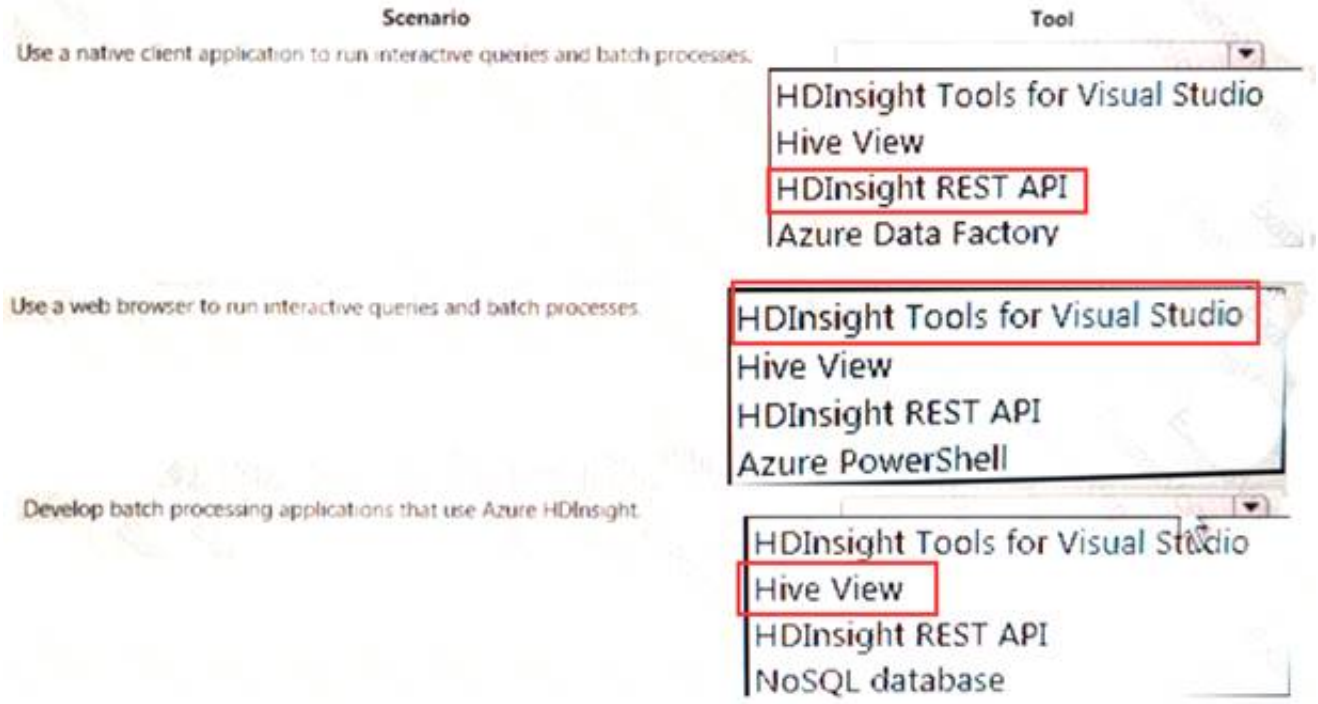

NEW QUESTION 20
- (Exam Topic 3)
A company has a real-lime data analysis solution that is hosted on Microsoft Azure the solution uses Azure Event Hub to ingest data and an Azure Stream Analytics cloud job to analyze the data. The cloud job is configured to use 120 Streaming Units (SU).
You need to optimize 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 23
- (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.

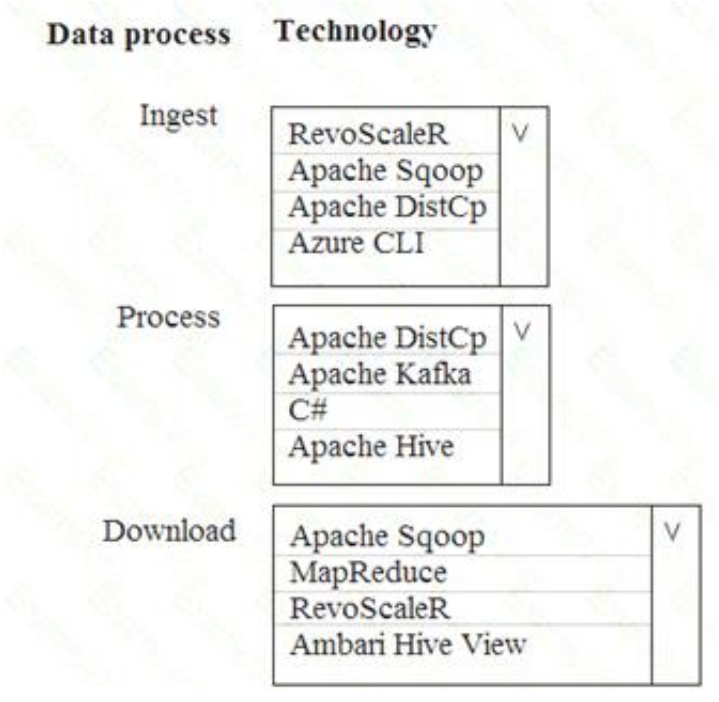

A. Mastered
B. Not Mastered

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 25
- (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 26
- (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.

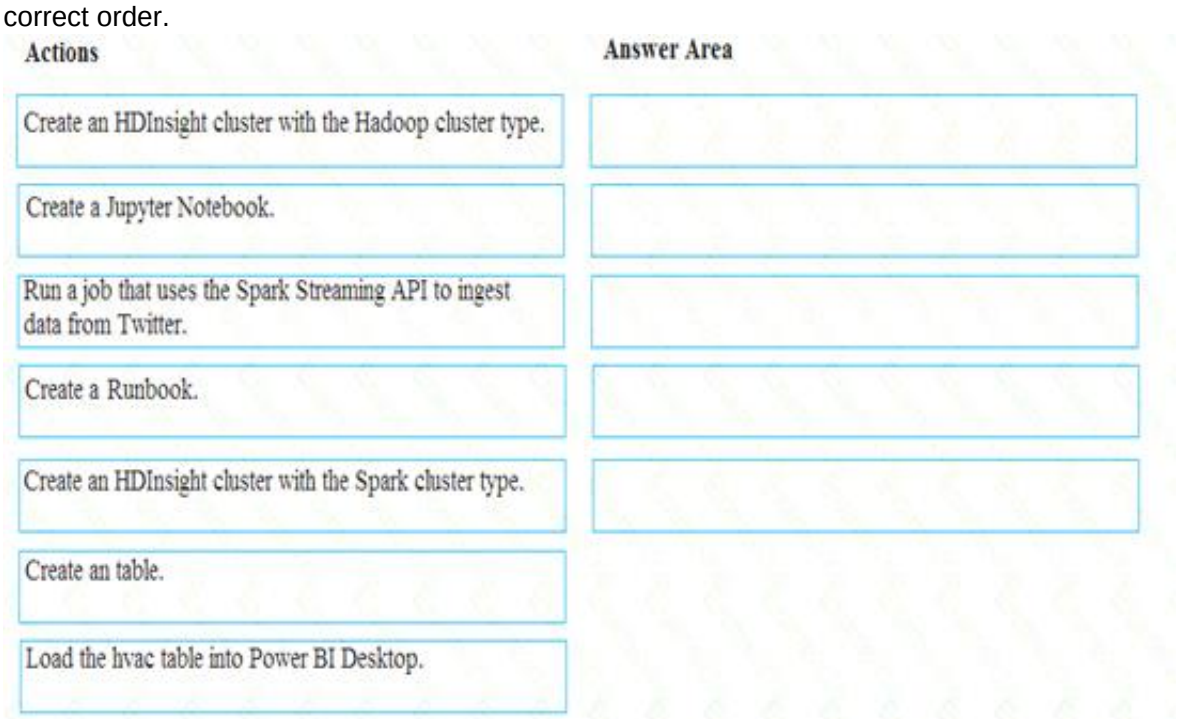

A. Mastered
B. Not Mastered

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 29
- (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 34
- (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

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

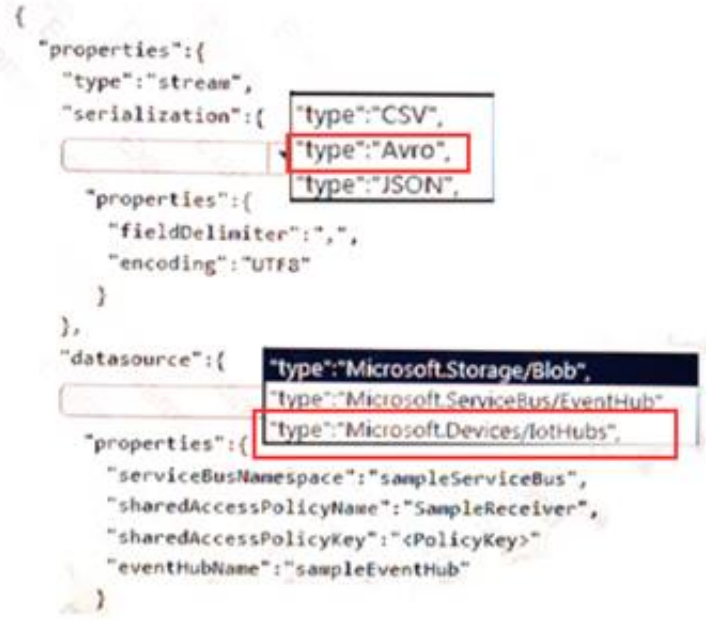

NEW QUESTION 35
- (Exam Topic 3)
You are developing a data engineering solution for a company. The solution will store a large set of key-value pair data by using Microsoft Azure Cosmos DB
The solution has the following requirements:
• Data must be partitioned into multiple containers.
• Data containers 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 39
- (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 43
- (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.

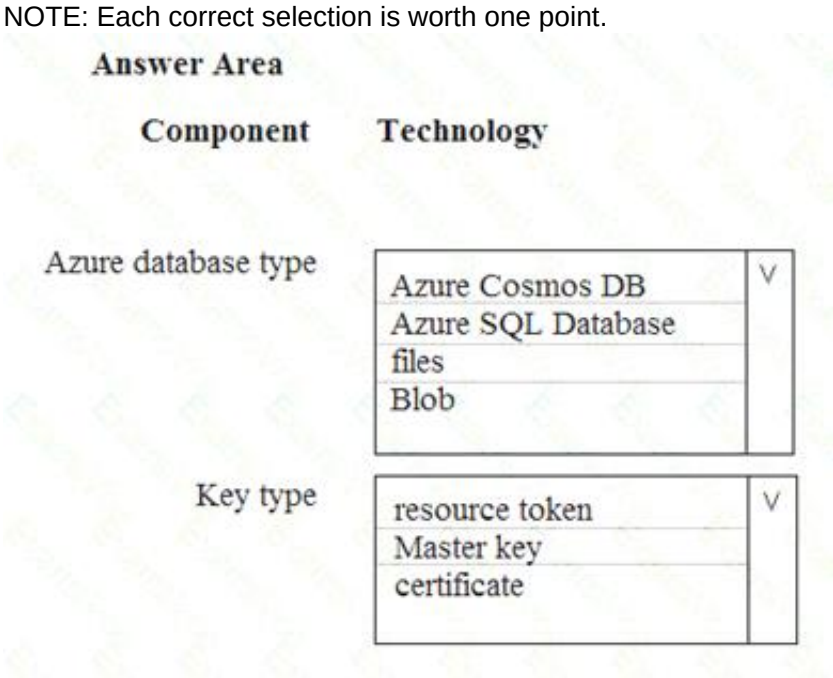

A. Mastered
B. Not Mastered

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 48
- (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 52
- (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 57
- (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.

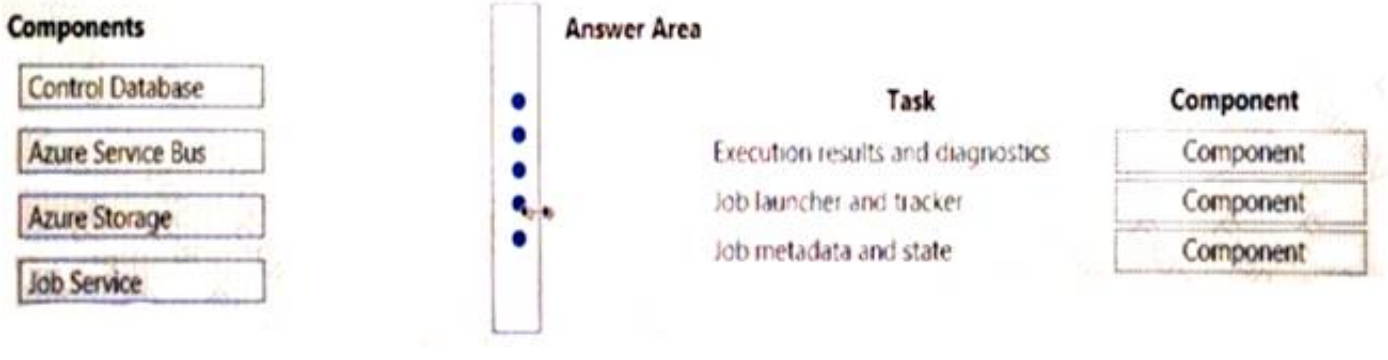

A. Mastered
B. Not Mastered

Answer: A

Explanation:

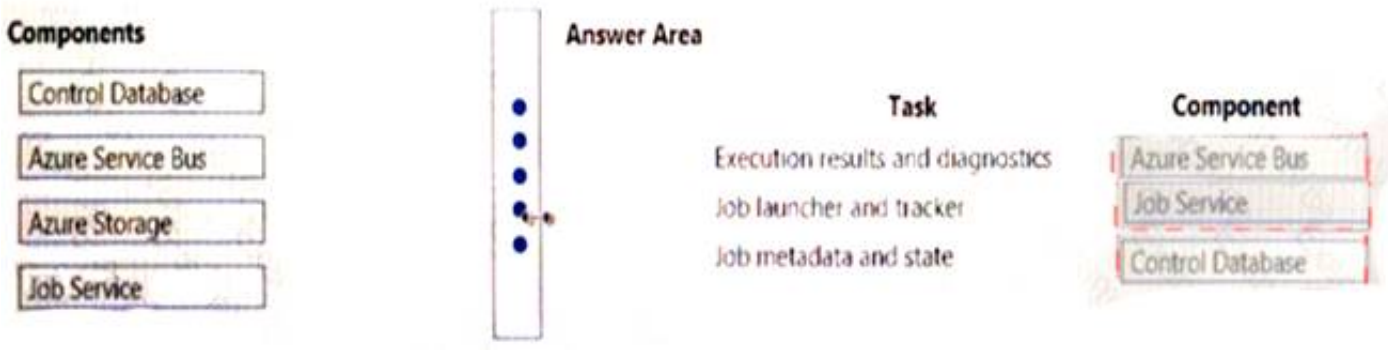

NEW QUESTION 58
- (Exam Topic 3)
A company builds an application to allow developers to share and compare code. The conversations, code snippets, and links shared by people in the application are stored in a Microsoft Azure SQL Database instance. The application allows for searches of historical conversations and code snippets.
When users share code snippets, the code snippet is compared against previously share 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.

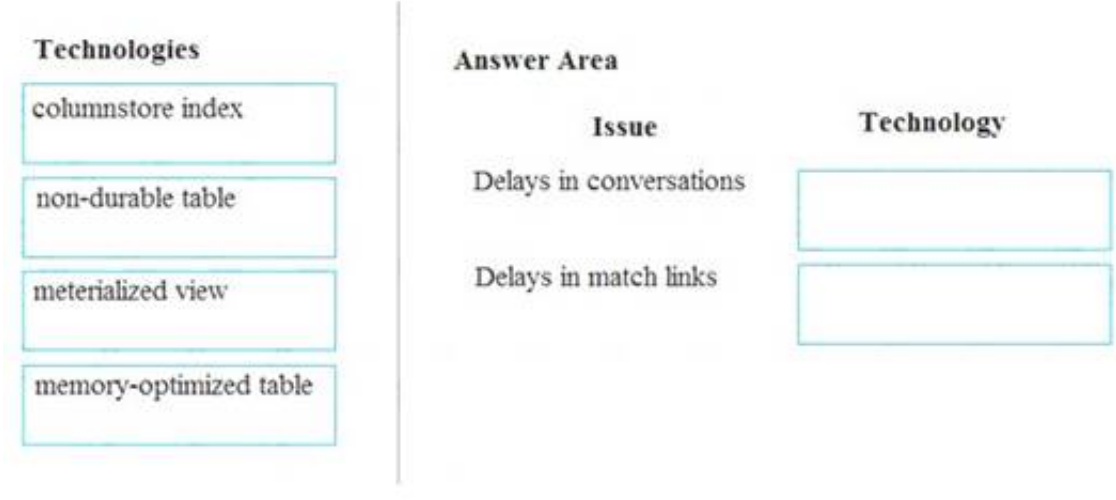

A. Mastered
B. Not Mastered

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 60
- (Exam Topic 3)
You develop data engineering solutions for a company.
You must integrate the company's on-premises Microsoft SQL Server data with Microsoft Azure SQL Database. Data must be transformed incrementally.
You need to implement the data integration solution.
Which tool should you use to configure a pipeline to copy data?

A. Use the Copy Data tool with Blob storage linked service as the source
B. Use Azure PowerShell with SQL Server linked service as a source
C. Use Azure Data Factory UI with Blob storage linked service as a source
D. Use the .NET Data Factory API with Blob storage linked service as the source

Answer: C

Explanation:
The Integration Runtime is a customer managed data integration infrastructure used by Azure Data Factory to provide data integration capabilities across different network environments.
A linked service defines the information needed for Azure Data Factory to connect to a data resource. We have three resources in this scenario for which linked services are needed:
- On-premises SQL Server
- Azure Blob Storage
- Azure SQL database

Note: Azure Data Factory is a fully managed cloud-based data integration service that orchestrates and automates the movement and transformation of data. The key concept in the ADF model is pipeline. A pipeline is a logical grouping of Activities, each of which defines the actions to perform on the data contained in Datasets. Linked services are used to define the information needed for Data Factory to connect to the data resources.
References:
https://docs.microsoft.com/en-us/azure/machine-learning/team-data-science-process/move-sql-azure-adf

NEW QUESTION 61
- (Exam Topic 3)
You are a data architect. The data engineering team needs to configure a synchronization of data between an on-premises Microsoft SQL Server database to Azure SQL Database.
Ad-hoc and reporting queries are being overutilized the on-premises production instance. The synchronization process must:
Perform an initial data synchronization to Azure SQL Database with minimal downtime Perform bi-directional data synchronization after initial synchronization
You need to implement this synchronization solution. Which synchronization method should you use?

A. transactional replication
B. Data Migration Assistant (DMA)
C. backup and restore

D. SQL Server Agent job
E. Azure SQL Data Sync

Answer: E

Explanation:
SQL Data Sync is a service built on Azure SQL Database that lets you synchronize the data you select bi-directionally across multiple SQL databases and SQL Server instances.
With Data Sync, you can keep data synchronized between your on-premises databases and Azure SQL databases to enable hybrid applications.
Compare Data Sync with Transactional Replication

	Data Sync	Transactional Replication
Advantages	- Active-active support - Bi-directional between on-premises and Azure SQL Database	- Lower latency - Transactional consistency - Reuse existing topology after migration
Disadvantages	- 5 min or more latency - No transactional consistency - Higher performance impact	- Can't publish from Azure SQL Database single database or pooled database - High maintenance cost

References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-sync-data

NEW QUESTION 65
- (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 70
- (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 74
- (Exam Topic 3)
A company has a SaaS solution that uses Azure SQL Database with elastic pools. The solution contains a dedicated database for each customer organization. Customer organizations have peak usage at different periods during the year.
You need to implement the Azure SQL Database elastic pool to minimize cost. Which option or options should you configure?

A. Number of transactions only
B. eDTUs per database only
C. Number of databases only
D. CPU usage only
E. eDTUs and max data size

Answer: E

Explanation:

The best size for a pool depends on the aggregate resources needed for all databases in the pool. This involves determining the following:
- Maximum resources utilized by all databases in the pool (either maximum DTUs or maximum vCores depending on your choice of resourcing model).
- Maximum storage bytes utilized by all databases in the pool.

Note: Elastic pools enable the developer to purchase resources for a pool shared by multiple databases to accommodate unpredictable periods of usage by individual databases. You can configure resources for the pool based either on the DTU-based purchasing model or the vCore-based purchasing model.
References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-pool

NEW QUESTION 75
- (Exam Topic 3)
A company has a SaaS solutions that will uses Azure SQL Database with elastic pools. The solution will have a dedicated database for each customer organization Customer organizations have peak usage at different periods during the year.
Which two factors affect your costs when sizing the Azure SQL Database elastic pools? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. maximum data size
B. number of databases
C. eDTUs consumption
D. number of read operations
E. number of transactions

Answer: AC

NEW QUESTION 79
- (Exam Topic 3)
You manage a Microsoft Azure SQL Data Warehouse Gen 2.
Users report slow performance when they run commonly used queries. Users do not report performance changes for infrequently used queries
You need to monitor resource utilization to determine the source of the performance issues. Which metric should you monitor?

A. Cache used percentage
B. Local tempdb percentage
C. WU percentage
D. CPU percentage

Answer: B

NEW QUESTION 81
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

DP-200 Practice Exam Features:

* DP-200 Questions and Answers Updated Frequently

* DP-200 Practice Questions Verified by Expert Senior Certified Staff

* DP-200 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* DP-200 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year