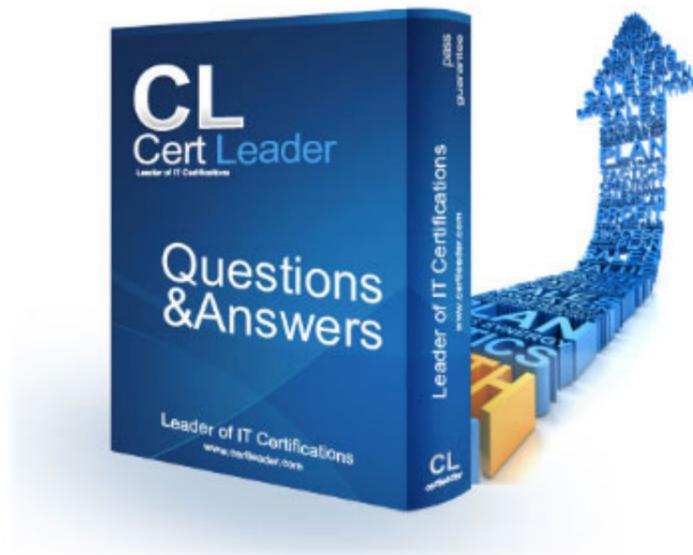


DP-200 Dumps

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 area 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 that phone-based polling 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 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	<table border="1"> <tr><td>FileCount</td><td></td></tr> <tr><td>BlobCapacity</td><td></td></tr> <tr><td>FileCapacity</td><td></td></tr> </table>	FileCount		BlobCapacity		FileCapacity	
FileCount							
BlobCapacity							
FileCapacity							
Aggregation	<table border="1"> <tr><td>Avg</td><td></td></tr> <tr><td>Sum</td><td></td></tr> </table>	Avg		Sum			
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)

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 5

- (Exam Topic 2)

You need set up the Azure Data Factory JSON definition for Tier 10 data.

What should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Data factory component	Value								
Connector	<table border="1"> <tr><td>connection string</td><td><input checked="" type="checkbox"/></td></tr> <tr><td>linked service name string</td><td><input type="checkbox"/></td></tr> <tr><td>gateway connection string</td><td><input type="checkbox"/></td></tr> <tr><td>data store name string</td><td><input type="checkbox"/></td></tr> </table>	connection string	<input checked="" type="checkbox"/>	linked service name string	<input type="checkbox"/>	gateway connection string	<input type="checkbox"/>	data store name string	<input type="checkbox"/>
connection string	<input checked="" type="checkbox"/>								
linked service name string	<input type="checkbox"/>								
gateway connection string	<input type="checkbox"/>								
data store name string	<input type="checkbox"/>								
Data movement activity	<table border="1"> <tr><td>Azure SQL Data Warehouse</td><td><input checked="" type="checkbox"/></td></tr> <tr><td>Azure Files</td><td><input type="checkbox"/></td></tr> <tr><td>Azure Blob</td><td><input type="checkbox"/></td></tr> <tr><td>Azure SQL Database</td><td><input type="checkbox"/></td></tr> </table>	Azure SQL Data Warehouse	<input checked="" type="checkbox"/>	Azure Files	<input type="checkbox"/>	Azure Blob	<input type="checkbox"/>	Azure SQL Database	<input type="checkbox"/>
Azure SQL Data Warehouse	<input checked="" type="checkbox"/>								
Azure Files	<input type="checkbox"/>								
Azure Blob	<input type="checkbox"/>								
Azure SQL Database	<input type="checkbox"/>								

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Connection String

To use storage account key authentication, you use the ConnectionString property, which specifies the information needed to connect to Blob Storage.

Mark this field as a SecureString to store it securely in Data Factory. You can also put account key in Azure Key Vault and pull the accountKey configuration out of the connection string.

Box 2: Azure Blob

Tier 10 reporting data must be stored in Azure Blobs

External Distribution and Sales	10	Yes, once ingested at Contoso main office	Data is ingested from multiple sources
---------------------------------	----	---	--

References:

<https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-blob-storage>

NEW QUESTION 6

- (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 the 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 7

- (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 8

- (Exam Topic 2)

You need to mask tier 1 data. Which functions should you use? To answer, select the appropriate option in the answer area.

NOTE: Each correct selection is worth one point.

Data type	Masking function								
A	<table border="1"> <tr><td>custom text</td><td><input type="checkbox"/></td></tr> <tr><td>default</td><td><input type="checkbox"/></td></tr> <tr><td>email</td><td><input type="checkbox"/></td></tr> <tr><td>random number</td><td><input type="checkbox"/></td></tr> </table>	custom text	<input type="checkbox"/>	default	<input type="checkbox"/>	email	<input type="checkbox"/>	random number	<input type="checkbox"/>
custom text	<input type="checkbox"/>								
default	<input type="checkbox"/>								
email	<input type="checkbox"/>								
random number	<input type="checkbox"/>								
B	<table border="1"> <tr><td>custom text</td><td><input type="checkbox"/></td></tr> <tr><td>default</td><td><input type="checkbox"/></td></tr> <tr><td>email</td><td><input type="checkbox"/></td></tr> <tr><td>random number</td><td><input type="checkbox"/></td></tr> </table>	custom text	<input type="checkbox"/>	default	<input type="checkbox"/>	email	<input type="checkbox"/>	random number	<input type="checkbox"/>
custom text	<input type="checkbox"/>								
default	<input type="checkbox"/>								
email	<input type="checkbox"/>								
random number	<input type="checkbox"/>								
C	<table border="1"> <tr><td>custom text</td><td><input type="checkbox"/></td></tr> <tr><td>default</td><td><input type="checkbox"/></td></tr> <tr><td>email</td><td><input type="checkbox"/></td></tr> <tr><td>random number</td><td><input type="checkbox"/></td></tr> </table>	custom text	<input type="checkbox"/>	default	<input type="checkbox"/>	email	<input type="checkbox"/>	random number	<input type="checkbox"/>
custom text	<input type="checkbox"/>								
default	<input type="checkbox"/>								
email	<input type="checkbox"/>								
random number	<input type="checkbox"/>								

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

A: Default

Full masking according to the data types of the designated fields.

For string data types, use XXXX or fewer Xs if the size of the field is less than 4 characters (char, nchar, varchar, nvarchar, text, ntext).

B: email

C: Custom text

Custom StringMasking method which exposes the first and last letters and adds a custom padding string in the middle. prefix,[padding],suffix

Tier 1 Database must implement data masking using the following masking logic:

Data type	Masking requirement
A	Mask 4 or less string data type characters
B	Mask first letter and domain
C	Mask everything except characters at the beginning and end

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking>

NEW QUESTION 9

- (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 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 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 and 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.

Azure SQL products	Database	Azure SQL product
Azure SQL Database elastic pools	Daily Sales	Azure SQL product
Azure SQL Database Premium	Weekly Analysis	Azure SQL product
Azure SQL Database Managed Instance		
Azure SQL Database Hyperscale		

- 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 17

- (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 Lake 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 19

- (Exam Topic 3)

You implement an event processing solution using Microsoft Azure Stream Analytics. The solution must meet the following requirements:

- Ingest data from Blob storage
- Analyze data in real time
- Store processed data in Azure Cosmos DB

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.

The screenshot shows an exam question interface with two main sections: 'Actions' and 'Answer Area'. The 'Actions' list contains six items:

- Create a query statement with the ORDER BY clause.
- Create a query statement with the SELECT INTO statement
- Configure Blob storage for a reference data JOIN clause
- Configure Azure Event Hub as input; select items with the TIMESTAMP BY clause.
- Set up Cosmos DB as the output
- Configure Blob storage as input, select items with the TIMESTAMP BY clause

There are two circular arrows (right and left) between the 'Actions' and 'Answer Area' sections, indicating that items can be moved between them. The 'Answer Area' is currently empty.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

The screenshot shows the same exam question interface as above, but with three actions moved from the 'Actions' list to the 'Answer Area' in the following sequence:

- Set up Cosmos DB as the output.
- Create a query statement with the SELECT INTO statement
- Configure Azure Event Hub as input; select items with the TIMESTAMP BY clause.

The 'Answer Area' also contains a right-pointing arrow icon, indicating the sequence is complete.

NEW QUESTION 22

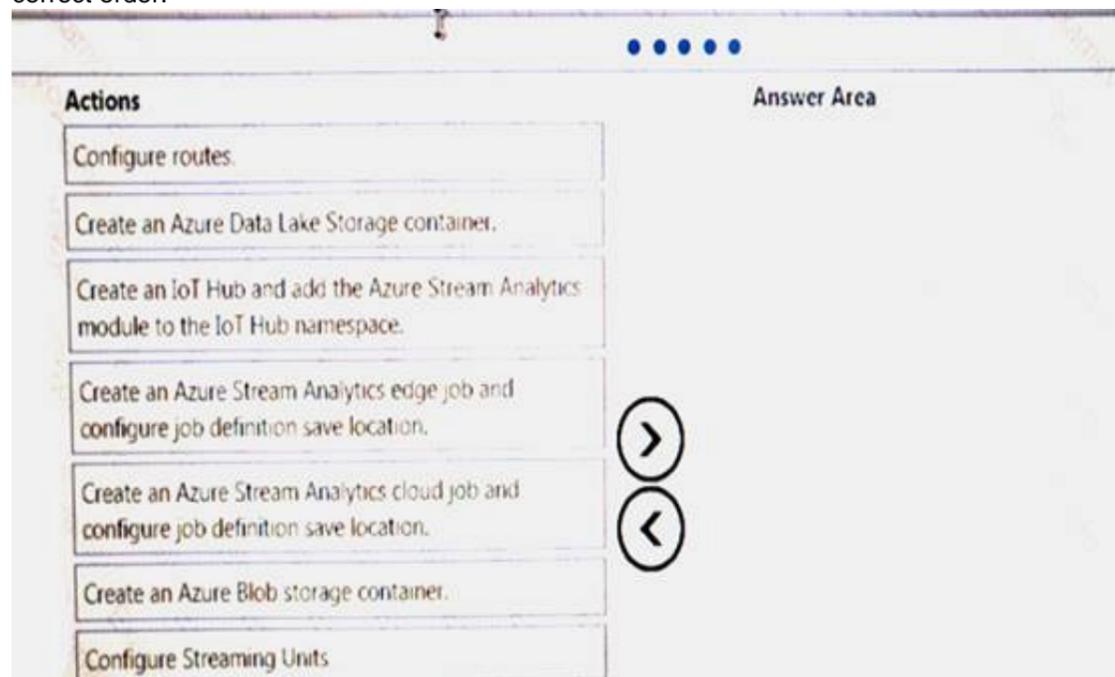
- (Exam Topic 3)

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.

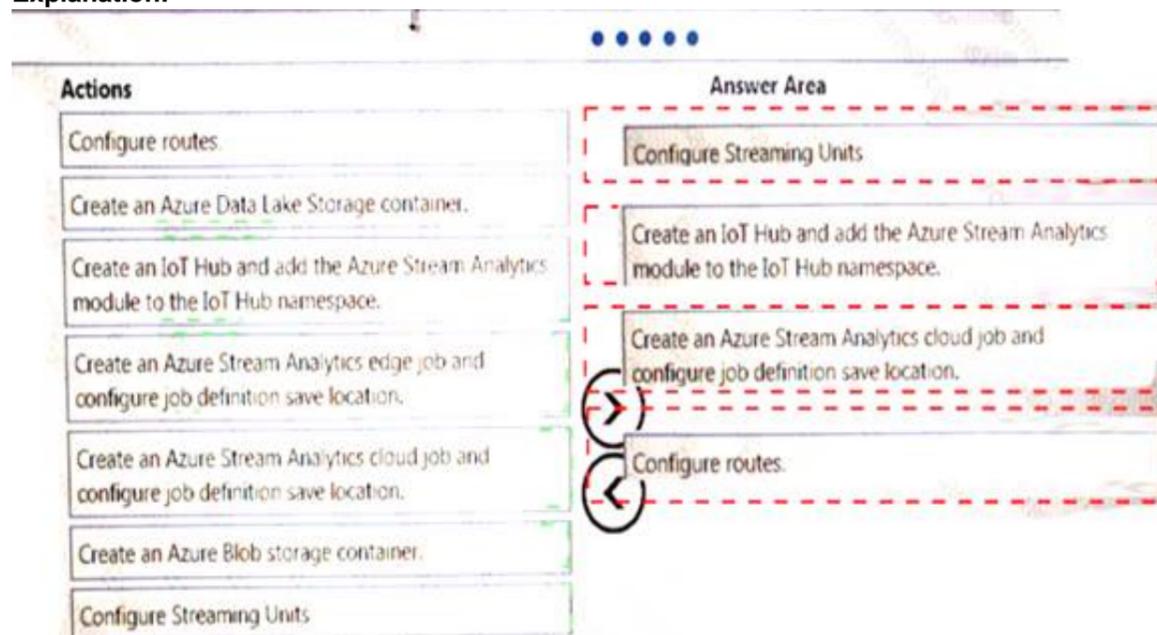
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 23

- (Exam Topic 3)

A company 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 26

- (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.

Data process	Technology								
Ingest	<table border="1"> <tr><td>RevoScaleR</td><td><input type="checkbox"/></td></tr> <tr><td>Apache Sqoop</td><td><input type="checkbox"/></td></tr> <tr><td>Apache DistCp</td><td><input type="checkbox"/></td></tr> <tr><td>Azure CLI</td><td><input type="checkbox"/></td></tr> </table>	RevoScaleR	<input type="checkbox"/>	Apache Sqoop	<input type="checkbox"/>	Apache DistCp	<input type="checkbox"/>	Azure CLI	<input type="checkbox"/>
RevoScaleR	<input type="checkbox"/>								
Apache Sqoop	<input type="checkbox"/>								
Apache DistCp	<input type="checkbox"/>								
Azure CLI	<input type="checkbox"/>								
Process	<table border="1"> <tr><td>Apache DistCp</td><td><input type="checkbox"/></td></tr> <tr><td>Apache Kafka</td><td><input type="checkbox"/></td></tr> <tr><td>C#</td><td><input type="checkbox"/></td></tr> <tr><td>Apache Hive</td><td><input type="checkbox"/></td></tr> </table>	Apache DistCp	<input type="checkbox"/>	Apache Kafka	<input type="checkbox"/>	C#	<input type="checkbox"/>	Apache Hive	<input type="checkbox"/>
Apache DistCp	<input type="checkbox"/>								
Apache Kafka	<input type="checkbox"/>								
C#	<input type="checkbox"/>								
Apache Hive	<input type="checkbox"/>								
Download	<table border="1"> <tr><td>Apache Sqoop</td><td><input type="checkbox"/></td></tr> <tr><td>MapReduce</td><td><input type="checkbox"/></td></tr> <tr><td>RevoScaleR</td><td><input type="checkbox"/></td></tr> <tr><td>Ambari Hive View</td><td><input type="checkbox"/></td></tr> </table>	Apache Sqoop	<input type="checkbox"/>	MapReduce	<input type="checkbox"/>	RevoScaleR	<input type="checkbox"/>	Ambari Hive View	<input type="checkbox"/>
Apache Sqoop	<input type="checkbox"/>								
MapReduce	<input type="checkbox"/>								
RevoScaleR	<input type="checkbox"/>								
Ambari Hive View	<input type="checkbox"/>								

- 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 28

- (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 31

- (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 32

- (Exam Topic 3)

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

You need to create the input stream.

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

NOTE: Each correct selection is worth one point.

Answer Area

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 35

- (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 36

- (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 to 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 41

- (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 44

- (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 48

- (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 49

- (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.

Azure PowerShell cmdlets	Answer Area
New-AzureRmSqlElasticPool	
New-AzureRmSqlServerFirewallRule	
New-AzureRmSqlServer	
New-AzureRmSqlServerVirtualNetworkRule	
New-AzureRmSqlDatabase	

- 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-powershell>

NEW QUESTION 54

- (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.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: memory-optimized table

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

Box 2: materialized view

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

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

References:

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

NEW QUESTION 58

- (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 59

- (Exam Topic 3)

A company plans to use Azure Storage for file storage purposes. Compliance rules require: A single storage account to store all operations including reads, writes and deletes

Retention of an on-premises copy of historical operations You need to configure the storage account.

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

- A. Configure the storage account to log read, write and delete operations for service type Blob
- B. Use the AzCopy tool to download log data from \$logs/blob
- C. Configure the storage account to log read, write and delete operations for service-type table
- D. Use the storage client to download log data from \$logs/table
- E. Configure the storage account to log read, write and delete operations for service type queue

Answer: AB

Explanation:

Storage Logging logs request data in a set of blobs in a blob container named \$logs in your storage account. This container does not show up if you list all the blob containers in your account but you can see its contents if you access it directly.

To view and analyze your log data, you should download the blobs that contain the log data you are interested in to a local machine. Many storage-browsing tools enable you to download blobs from your storage account; you can also use the Azure Storage team provided command-line Azure Copy Tool (AzCopy) to download your log data.

References:

<https://docs.microsoft.com/en-us/rest/api/storageservices/enabling-storage-logging-and-accessing-log-data>

NEW QUESTION 61

- (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 65

- (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 66

- (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 load 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 70

- (Exam Topic 3)

You are the data engineer for your company. An application uses a NoSQL database to store data. The database uses the key-value and wide-column NoSQL database type.

Developers need to access data in the database using an API.

You need to determine which API to use for the database model and type.

Which two APIs should you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Table API
- B. MongoDB API
- C. Gremlin API
- D. SQL API
- E. Cassandra API

Answer: BE

Explanation:

B: Azure Cosmos DB is the globally distributed, multimodel database service from Microsoft for mission-critical applications. It is a multimodel database and supports document, key-value, graph, and columnar data models.

E: Wide-column stores store data together as columns instead of rows and are optimized for queries over large datasets. The most popular are Cassandra and HBase.

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/graph-introduction> <https://www.mongodb.com/scale/types-of-nosql-databases>

NEW QUESTION 75

- (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 77

- (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 78

- (Exam Topic 3)

You are a data engineer. You are designing a Hadoop Distributed File System (HDFS) architecture. You plan to use Microsoft Azure Data Lake as a data storage repository.

You must provision the repository with a resilient data schema. You need to ensure the resiliency of the Azure Data Lake Storage. What should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Requirement	Node				
Provide data access to clients.	<table border="1"> <tr> <td>DataNode</td> <td><input type="checkbox"/></td> </tr> <tr> <td>NameNode</td> <td><input type="checkbox"/></td> </tr> </table>	DataNode	<input type="checkbox"/>	NameNode	<input type="checkbox"/>
DataNode	<input type="checkbox"/>				
NameNode	<input type="checkbox"/>				
Run operations on files and directories of the file system.	<table border="1"> <tr> <td>DataNode</td> <td><input type="checkbox"/></td> </tr> <tr> <td>NameNode</td> <td><input type="checkbox"/></td> </tr> </table>	DataNode	<input type="checkbox"/>	NameNode	<input type="checkbox"/>
DataNode	<input type="checkbox"/>				
NameNode	<input type="checkbox"/>				
Perform block creation, deletion, and replication.	<table border="1"> <tr> <td>DataNode</td> <td><input type="checkbox"/></td> </tr> <tr> <td>NameNode</td> <td><input type="checkbox"/></td> </tr> </table>	DataNode	<input type="checkbox"/>	NameNode	<input type="checkbox"/>
DataNode	<input type="checkbox"/>				
NameNode	<input type="checkbox"/>				

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: NameNode

An HDFS cluster consists of a single NameNode, a master server that manages the file system namespace and regulates access to files by clients.

Box 2: DataNode

The DataNodes are responsible for serving read and write requests from the file system's clients. Box 3: DataNode
The DataNodes perform block creation, deletion, and replication upon instruction from the NameNode.

Note: HDFS has a master/slave architecture. An HDFS cluster consists of a single NameNode, a master server that manages the file system namespace and regulates access to files by clients. In addition, there are a number of DataNodes, usually one per node in the cluster, which manage storage attached to the nodes that they run on. HDFS exposes a file system namespace and allows user data to be stored in files. Internally, a file is split into one or more blocks and these blocks are stored in a set of DataNodes. The NameNode executes file system namespace operations like opening, closing, and renaming files and directories. It also determines the mapping of blocks to DataNodes. The DataNodes are responsible for serving read and write requests from the file system's clients. The DataNodes also perform block creation, deletion, and replication upon instruction from the NameNode.

References: https://hadoop.apache.org/docs/r1.2.1/hdfs_design.html#NameNode+and+DataNodes

NEW QUESTION 80

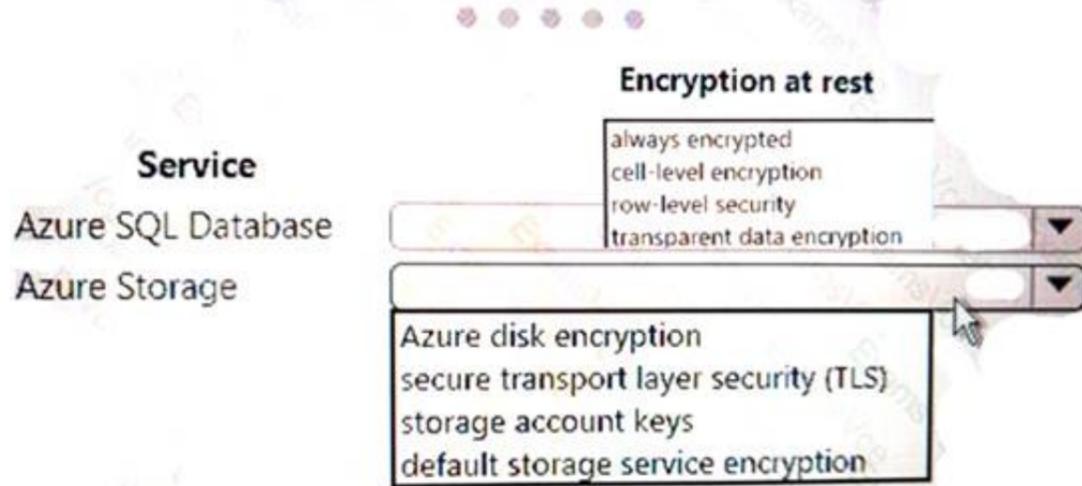
- (Exam Topic 3)

Your company uses Azure SQL Database and Azure Blob storage.

All data at rest must be encrypted by using the company's own key. The solution must minimize administrative effort and the impact to applications which use the database.

You need to configure security.

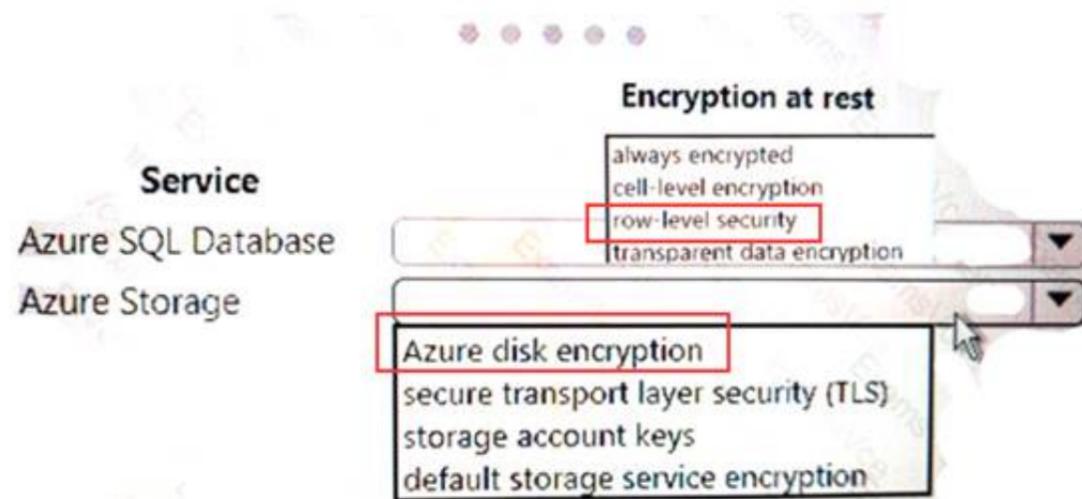
What should you implement? To answer, select the appropriate option in the answer area. NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 83

- (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 88

.....

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>