Exam Questions DP-200

Implementing an Azure Data Solution

https://www.2passeasy.com/dumps/DP-200/

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

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

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

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

NEW QUESTION 2
- (Exam Topic 1)
You need to ensure that Azure Data Factory pipelines can be deployed. How should you configure authentication and authorization for deployments? To answer, select the appropriate options in the answer choices.
NOTE: Each correct selection is worth one point.

Security requirement	Technology
Authorization	RBAC ∨ DAC MAC Claims
Authentication	Service Principal ∧ Kerberos Certificate-based Bearer Token

A. Mastered
B. Not Mastered

Answer: A

Explanation:
The way you control access to resources using RBAC is to create role assignments. This is a key concept to understand – it's how permissions are enforced. A

role assignment consists of three elements: security principal, role definition, and scope.
Scenario:
No credentials or secrets should be used during deployments
Phone-based poll data must only be uploaded by authorized users from authorized devices Contractors must not have access to any polling data other than their own
Access to polling data must set on a per-active directory user basis References:
https://docs.microsoft.com/en-us/azure/role-based-access-control/overview

NEW QUESTION 3
- (Exam Topic 1)
You need to ensure polling data security requirements are met.
Which security technologies should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Azure Active Directory user Scenario:
Access to polling data must set on a per-active directory user basis
Box 2: DataBase Scoped Credential
SQL Server uses a database scoped credential to access non-public Azure blob storage or Kerberos-secured Hadoop clusters with PolyBase.
PolyBase cannot authenticate by using Azure AD authentication. References:
https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-scoped-credential-transact-sql

NEW QUESTION 4
- (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	custom text	∨
	default	
	email	
	random number	
B	custom text	∨
	default	
	email	
	random number	
C	custom text	∨
	default	
	email	
	random number	

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 5
- (Exam Topic 3)
You are designing a new Lambda architecture on Microsoft Azure. The real-time processing layer must meet the following requirements: Ingestion:
- Receive millions of events per second
- Act as a fully managed Platform-as-a-Service (PaaS) solution
- Integrate with Azure Functions

Stream processing:
- Process on a per-job basis
- Provide seamless connectivity with Azure services
- Use a SQL-based query language

Analytical data store:
- Act as a managed service
- Use a document store
- Provide data encryption at rest

You need to identify the correct technologies to build the Lambda architecture using minimal effort. Which technologies should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Architecture requirement / **Answer Area**

Ingestion:
- HDInsight Kafka
- Azure Event Hubs
- HDInsight Storm
- HDInsight Spark

Stream Processing:
- Azure Stream Analytics
- HDInsight with Spark Streaming
- Azure Cosmos DB Change Feed
- Azure Analysis Services

Analytical Data Store:
- Hive LLAP on HDInsight
- Azure Analysis Services
- Azure Cosmos DB
- SQL Data Warehouse

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Azure Event Hubs
This portion of a streaming architecture is often referred to as stream buffering. Options include Azure Event Hubs, Azure IoT Hub, and Kafka.

NEW QUESTION 6
- (Exam Topic 3)
Your company manages on-premises Microsoft SQL Server pipelines by using a custom solution.
The data engineering team must implement a process to pull data from SQL Server and migrate it to Azure Blob storage. The process must orchestrate and manage the data lifecycle.
You need to configure Azure Data Factory to connect to the on-premises SQL Server database.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions / **Answer Area**

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Create a virtual private network (VPN) connection from on-premises to Microsoft Azure.
You can also use IPSec VPN or Azure ExpressRoute to further secure the communication channel between your on-premises network and Azure.
Azure Virtual Network is a logical representation of your network in the cloud. You can connect an
on-premises network to your virtual network by setting up IPSec VPN (site-to-site) or ExpressRoute (private peering).
Step 2: Create an Azure Data Factory resource. Step 3: Configure a self-hosted integration runtime.
You create a self-hosted integration runtime and associate it with an on-premises machine with the SQL Server database. The self-hosted integration runtime is the component that copies data from the SQL Server database on your machine to Azure Blob storage.
Note: A self-hosted integration runtime can run copy activities between a cloud data store and a data store in a private network, and it can dispatch transform activities against compute resources in an on-premises network or an Azure virtual network. The installation of a self-hosted integration runtime needs on an on-premises machine or a virtual machine (VM) inside a private network.
References:
https://docs.microsoft.com/en-us/azure/data-factory/tutorial-hybrid-copy-powershell

NEW QUESTION 7
- (Exam Topic 3)
You develop data engineering solutions for a company.
A project requires an in-memory batch data processing solution.
You need to provision an HDInsight cluster for batch processing of data on Microsoft Azure.
How should you complete the PowerShell segment? To answer, select the appropriate option in the answer area.
NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 8
- (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 9
- (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.

Actions

| In Auditing, enable **Auditing.** |
| Configure the service to create alerts for threat detections of type **Data Exfiltration.** |
| In Firewalls and virtual networks, enable **Allow access to Azure services.** |
| Enable advanced threat protection. |
| Configure the service to send email alerts to security@contoso.com |

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Actions

| In Auditing, enable **Auditing.** |
| Configure the service to create alerts for threat detections of type **Data Exfiltration.** |
| In Firewalls and virtual networks, enable **Allow access to Azure services.** |
| Enable advanced threat protection. |
| Configure the service to send email alerts to security@contoso.com |

Answer Area

| Enable advanced threat protection. |
| Configure the service to send email alerts to security@contoso.com |
| Configure the service to create alerts for threat detections of type **Data Exfiltration.** |

NEW QUESTION 10
- (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 10
- (Exam Topic 3)

A company plans to use Azure SQL Database to support a mission-critical application.
The application must be highly available without performance degradation during maintenance windows. You need to implement the solution.
Which three technologies should you implement? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Premium service tier
B. Virtual machine Scale Sets
C. Basic service tier
D. SQL Data Sync
E. Always On availability groups
F. Zone-redundant configuration

Answer: AEF

Explanation:
Premium/business critical service tier model that is based on a cluster of database engine processes. This architectural model relies on a fact that there is always a quorum of available database engine nodes and has minimal performance impact on your workload even during maintenance activities.
In the premium model, Azure SQL database integrates compute and storage on the single node. High availability in this architectural model is achieved by replication of compute (SQL Server Database Engine process) and storage (locally attached SSD) deployed in 4-node cluster, using technology similar to SQL Server Always On Availability Groups.
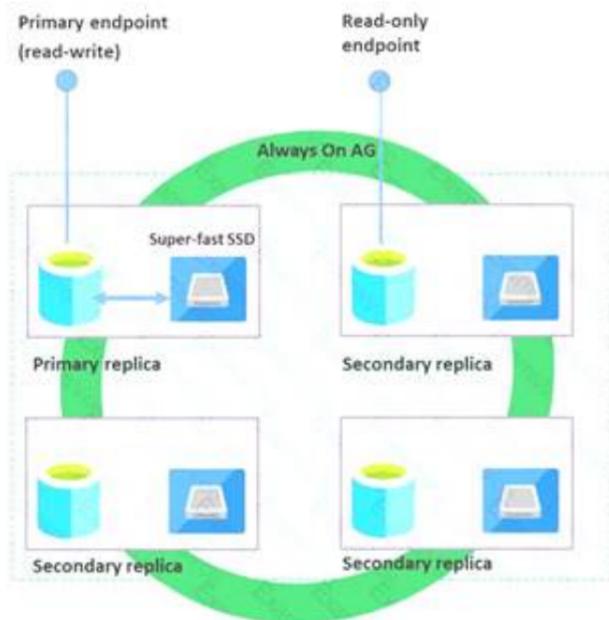
Business Critical service tier: collocated compute and storage

Zone redundant configuration
By default, the quorum-set replicas for the local storage configurations are created in the same datacenter. With the introduction of Azure Availability Zones, you have the ability to place the different replicas in the quorum-sets to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).
References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-high-availability

NEW QUESTION 15
- (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
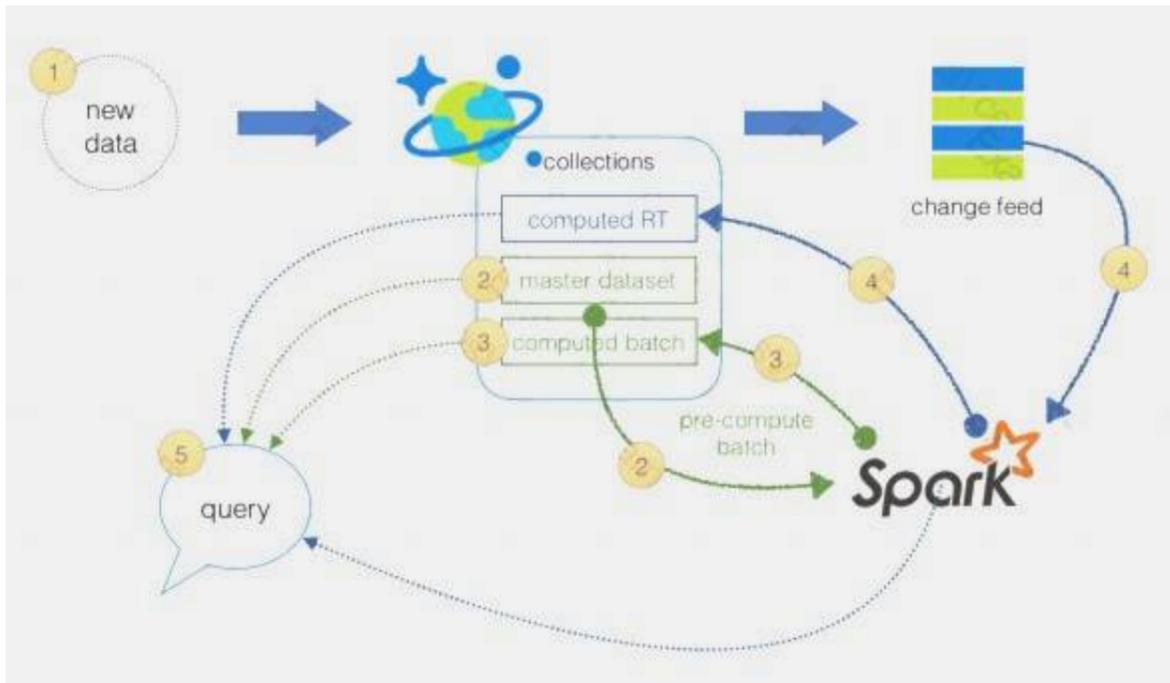
Note: Lambda architecture is a data-processing architecture designed to handle massive quantities of data by taking advantage of both batch processing and stream processing methods, and minimizing the latency involved in querying big data.
References:
https://sqlwithmanoj.com/2018/02/16/what-is-lambda-architecture-and-what-azure-offers-with-its-new-cosmos-

NEW QUESTION 18
- (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 21

- (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 24
- (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 28
- (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 31
- (Exam Topic 3)
Note: This question is part of series of questions that present the same scenario. Each question in the series contain a unique solution. Determine whether the solution meets the stated goals.
You develop data engineering solutions for a company.
A project requires the deployment of resources to Microsoft Azure for batch data processing on Azure HDInsight. Batch processing will run daily and must: Scale to minimize costs
Be monitored for cluster performance
You need to recommend a tool that will monitor clusters and provide information to suggest how to scale. Solution: Monitor cluster load using the Ambari Web UI. Does the solution meet the goal?

A. Yes
B. No

Answer: B

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

NEW QUESTION 33
- (Exam Topic 3)
A company manages several on-premises Microsoft SQL Server databases.
You need to migrate the databases to Microsoft Azure by using a backup and restore process. Which data technology should you use?

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

Answer: D

Explanation:
Managed instance is a new deployment option of Azure SQL Database, providing near 100% compatibility with the latest SQL Server on-premises (Enterprise Edition) Database Engine, providing a native virtual network (VNet) implementation that addresses common security concerns, and a business model favorable for on-premises SQL Server customers. The managed instance deployment model allows existing SQL Server customers to lift and shift their on-premises applications to the cloud with minimal application and database changes.
References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-managed-instance

NEW QUESTION 35
- (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:

Components		Answer Area		
Control Database			**Task**	**Component**
Azure Service Bus		•	Execution results and diagnostics	Azure Service Bus
Azure Storage		•	Job launcher and tracker	Job Service
Job Service		• •	Job metadata and state	Control Database
		•		
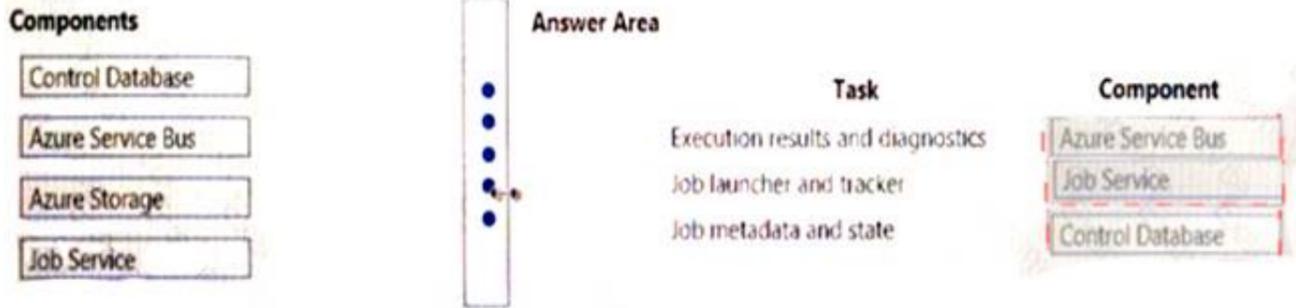
NEW QUESTION 36
- (Exam Topic 3)
You plan to create a new single database instance of Microsoft Azure SQL Database.
The database must only allow communication from the data engineer's workstation. You must connect directly to the instance by using Microsoft SQL Server Management Studio.
You need to create 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-powersh

NEW QUESTION 39
- (Exam Topic 3)
A company runs Microsoft Dynamics CRM with Microsoft SQL Server on-premises. SQL Server Integration Services (SSIS) packages extract data from Dynamics CRM APIs, and load the data into a SQL Server data warehouse.
The datacenter is running out of capacity. Because of the network configuration, you must extract on premises data to the cloud over https. You cannot open any additional ports. The solution must implement the least amount of effort.
You need to create the pipeline system.
Which component should you use? To answer, select the appropriate technology in the dialog box in the answer area.
NOTE: Each correct selection is worth one point.

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

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 41
- (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 44
- (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 48
- (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.

Actions	Answer Area
Upload a certificate	
Add secrets to the scope	
Use Blob Storage access key	
Create a secret scope	
Configure a JDBC connector	
Mount the Azure Blob Storage container	

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 53
- (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 57
- (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 58
- (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 60
- (Exam Topic 3)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
A company uses Azure Data Lake Gen 1 Storage to store big data related to consumer behavior. You need to implement logging.
Solution: Use information stored m Azure Active Directory reports.
Does the solution meet the goal?

A. Yes
B. No

Answer: B

NEW QUESTION 63
- (Exam Topic 3)
You implement 3 Azure SQL Data Warehouse instance.
You plan to migrate the largest fact table to Azure SQL Data Warehouse The table resides on Microsoft SQL Server on-premises and e 10 terabytes (TB) in size.
Incoming queues use the primary key Sale Key column to retrieve data as displayed in the following table:

SaleKey	CityKey	CustomerKey	StockItemKey	InvoiceDateKey	Quantity	UnitPrice	TotalExcludingTax
49309	90852	70	69	10/22/13	8	16	128
49313	55710	126	69	10/22/13	2	16	32
49343	44710	234	68	10/22/13	10	16	160
49352	66109	163	70	10/22/13	4	16	64
49448	65312	230	70	10/22/13	8	16	128
49646	85877	271	70	10/24/13	1	16	16
49798	41238	288	69	10/24/13	1	16	16

You need to distribute the fact table across multiple nodes to optimize performance of the table. Which technology should you use?

A. hash distributed table with clustered ColumnStore index
B. hash distributed table with clustered index
C. heap table with distribution replicate
D. round robin distributed table with clustered index
E. round robin distributed table with clustered ColumnStore index

Answer: A

NEW QUESTION 65
- (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 70
- (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.	DataNode / NameNode [V]
Run operations on files and directories of the file system.	DataNode / NameNode [V]
Perform block creation, deletion, and replication.	DataNode / NameNode [V]

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 75
- (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 76
- (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 79
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DP-200 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DP-200 Product From:

https://www.2passeasy.com/dumps/DP-200/

Money Back Guarantee

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