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 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 2
- (Exam Topic 1)
You need to ensure phone-based polling data upload reliability requirements are met. How should you configure monitoring? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Setting	Value
Metric	FileCount BlobCapacity FileCapacity
Aggregation	Avg Sum

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 3
- (Exam Topic 2)
You need to set up access to Azure SQL Database for Tier 7 and Tier 8 partners.
Which three actions should you perform in sequence? To answer, move the appropriate three actions from the list of actions to the answer area and arrange them in the correct order.

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Tier 7 and 8 data access is constrained to single endpoints managed by partners for access Step 1: Set the Allow Azure Services to Access Server setting to Disabled

Set Allow access to Azure services to OFF for the most secure configuration.

By default, access through the SQL Database firewall is enabled for all Azure services, under Allow access to Azure services. Choose OFF to disable access for all Azure services.

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

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

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

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

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

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

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

To setup a database-level firewall rule:

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

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

☉ On the toolbar, select Execute to create the firewall rule. References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-security-tutorial

NEW QUESTION 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)
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
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 6
- (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	Answer Area
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	

A. Mastered
B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 7
- (Exam Topic 3)
You are developing the data platform for a global retail company. The company operates during normal working hours in each region. The analytical database is used once a week for building sales projections.
Each region maintains its own private virtual network.
Building the sales projections is very resource intensive are generates upwards of 20 terabytes (TB) of data. Microsoft Azure SQL Databases must be provisioned.
🔘 Database provisioning must maximize performance and minimize cost
🔘 The daily sales for each region must be stored in an Azure SQL Database instance
🔘 Once a day, the data for all regions must be loaded in an analytical Azure SQL Database instance You need to provision Azure SQL database instances.
How should you provision the database instances? To answer, drag the appropriate Azure SQL products to the correct databases. Each Azure SQL product may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

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 8
- (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.
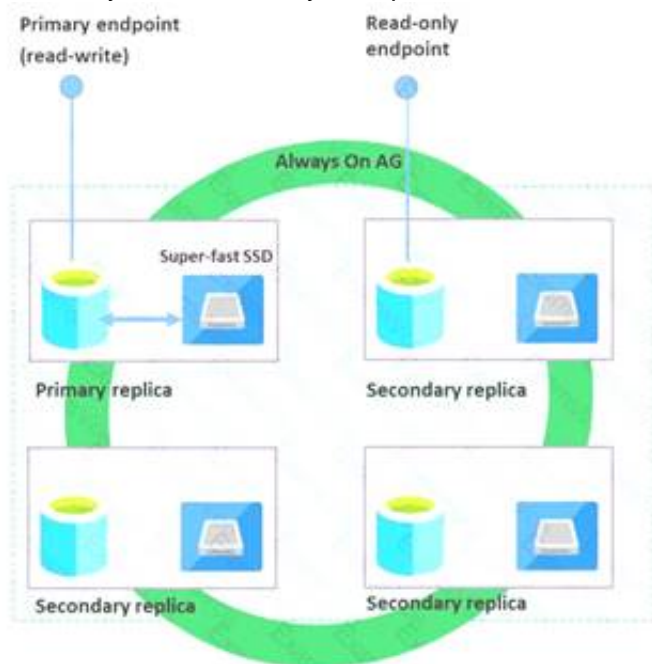NOTE: Each correct selection is worth one point.

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 9
- (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.
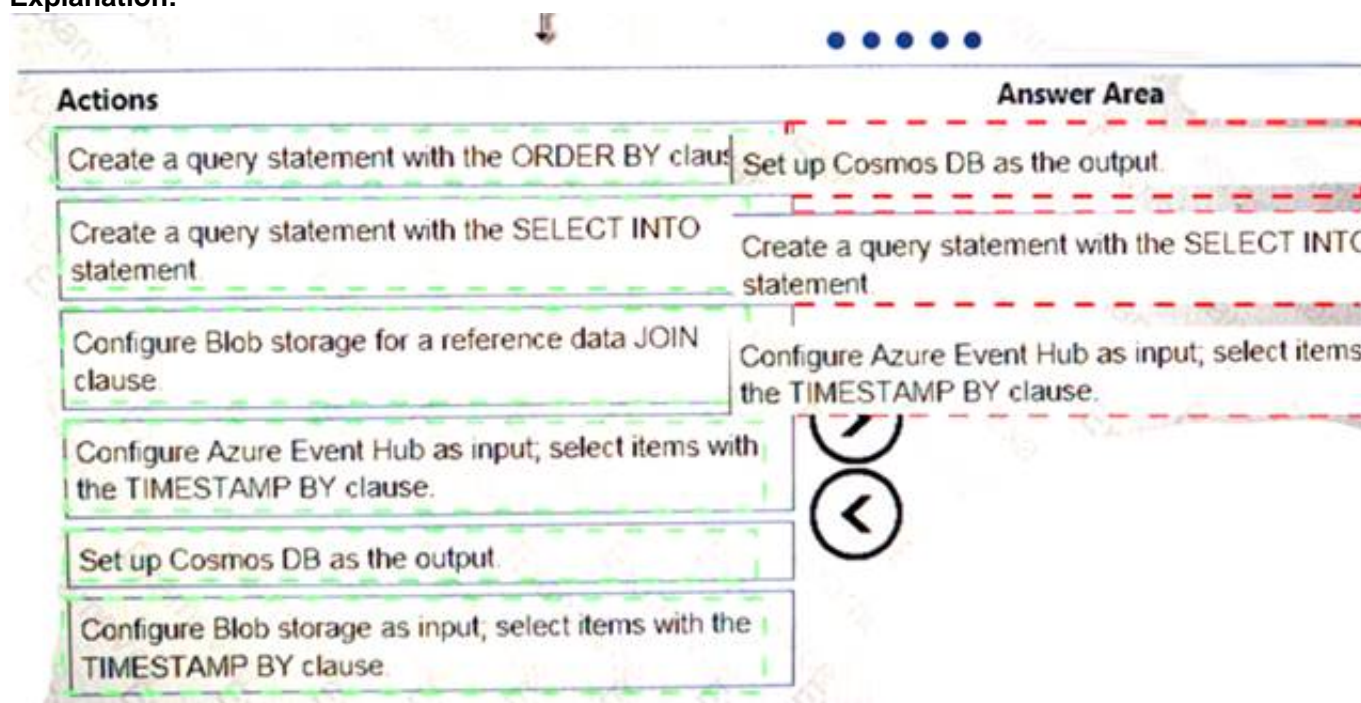

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 10
- (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 10
- (Exam Topic 3)
A company is deploying a service-based data environment. You are developing a solution to process this data. The solution must meet the following requirements:
🔘 Use an Azure HDInsight cluster for data ingestion from a relational database in a different cloud service
🔘 Use an Azure Data Lake Storage account to store processed data
🔘 Allow users to download processed data
You need to recommend technologies for the solution.
Which technologies should you use? To answer, select the appropriate options in the answer area.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Apache Sqoop is a tool designed for efficiently transferring bulk data between Apache Hadoop and structured datastores such as relational databases.
Azure HDInsight is a cloud distribution of the Hadoop components from the Hortonworks Data Platform (HDP).

NEW QUESTION 13
- (Exam Topic 3)
You develop data engineering solutions for a company.
A project requires analysis of real-time Twitter feeds. Posts that contain specific keywords must be stored and processed on Microsoft Azure and then displayed by using Microsoft Power BI. You need to implement the solution.
Which five actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

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

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 18
- (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 21
- (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 24
- (Exam Topic 3)
You plan to create a new single database instance of Microsoft Azure SQL Database.
The database must only allow communication from the data engineer's workstation. You must connect directly to the instance by using Microsoft SQL Server Management Studio.
You need to create and configure the Database. Which three Azure PowerShell cmdlets should you use to develop the solution? To answer, move the appropriate cmdlets from the list of cmdlets to the answer area and arrange them in the correct order.

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 29
- (Exam Topic 3)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.
You develop a data ingestion process that will import data to a Microsoft Azure SQL Data Warehouse.
The data to be ingested resides in parquet files stored in an Azure Data lake Gen 2 storage account.
You need to load the data from the Azure Data Lake Gen 2 storage account into the Azure SQL Data Warehouse.
Solution;
1. Create an external data source pointing to the Azure Data Lake Gen 2 storage account.
2. Create an external tile format and external table using the external data source.
3. Load the data using the CREATE TABLE AS SELECT statement.
Does the solution meet the goal?

A. Yes
B. No

Answer: A

NEW QUESTION 31
- (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 36
- (Exam Topic 3)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.
You develop a data ingestion process that will import data to a Microsoft Azure SQL Data Warehouse. The data to be ingested resides in parquet files stored in an Azure Data lake Gen 2 storage account.
You need to load the data from the Azure Data Lake Gen 2 storage account into the Azure SQL Data Warehouse.
Solution:
1. Create an external data source pointing to the Azure storage account
2. Create a workload group using the Azure storage account name as the pool name
3. Load the data using the INSERT…SELECT statement
Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
You need to create an external file format and external table using the external data source. You then load the data using the CREATE TABLE AS SELECT statement.
References:
https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-load-from-azure-data-lake-store
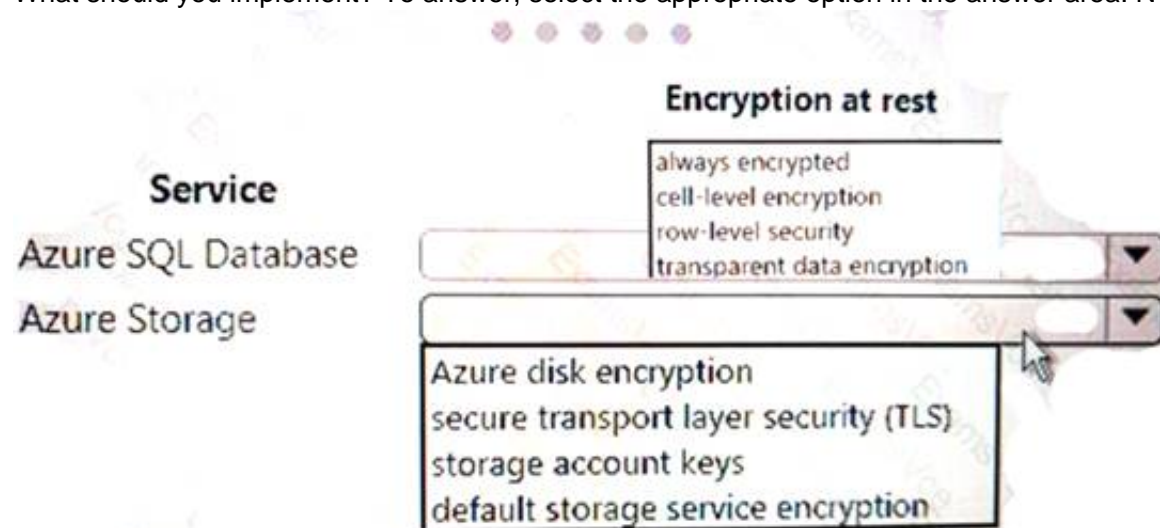
NEW QUESTION 41
- (Exam Topic 3)
Your company uses Azure SQL Database and Azure Blob storage.
All data at rest must be encrypted by using the company's own key. The solution must minimize administrative effort and the impact to applications which use the database.
You need to configure security.
What 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:

Encryption at rest

Service	
Azure SQL Database	always encrypted cell-level encryption row-level security transparent data encryption
Azure Storage	Azure disk encryption secure transport layer security (TLS) storage account keys default storage service encryption

NEW QUESTION 45
......

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