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

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

Answer: EF

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

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

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

NEW QUESTION 4
- (Exam Topic 3)
Contoso, Ltd. plans to configure existing applications to use Azure SQL Database. When security-related operations occur, the security team must be informed.
You need to configure Azure Monitor while minimizing administrative efforts
Which three actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Create a new action group to email alerts@contoso.com.
B. Use alerts@contoso.com as an alert email address.
C. Use all security operations as a condition.
D. Use all Azure SQL Database servers as a resource.
E. Query audit log entries as a condition.

Answer: ACE

NEW QUESTION 5
- (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 6
- (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 7
- (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 8
- (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.

Actions

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Actions

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

Answer Area

Set up Cosmos DB as the output.

Create a query statement with the SELECT INTO statement

Configure Azure Event Hub as input; select items the TIMESTAMP BY clause.

NEW QUESTION 9
- (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
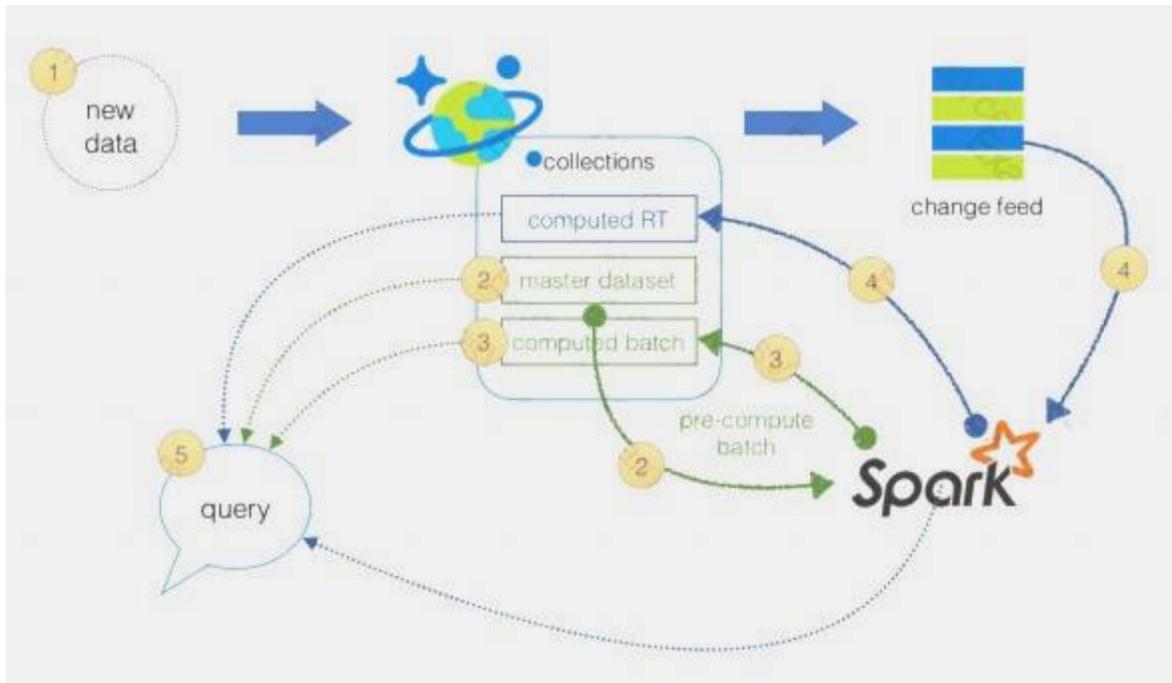
Note: Lambda architecture is a data-processing architecture designed to handle massive quantities of data by taking advantage of both batch processing and stream processing methods, and minimizing the latency involved in querying big data.
References:
https://sqlwithmanoj.com/2018/02/16/what-is-lambda-architecture-and-what-azure-offers-with-its-new-cosmos-

NEW QUESTION 10
- (Exam Topic 3)
A company has a real-lime data analysis solution that is hosted on Microsoft Azure the solution uses Azure Event Hub to ingest data and an Azure Stream Analytics cloud job to analyze the data. The cloud job is configured to use 120 Streaming Units (SU).
You need to optimize performance for the Azure Stream Analytics job.
Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one port.

A. Implement event ordering
B. Scale the SU count for the job up
C. Implement Azure Stream Analytics user-defined functions (UDF)
D. Scale the SU count for the job down
E. Implement query parallelization by partitioning the data output
F. Implement query parallelization by partitioning the data input

Answer: BF

Explanation:
Scale out the query by allowing the system to process each input partition separately.
F: A Stream Analytics job definition includes inputs, a query, and output. Inputs are where the job reads the data stream from.
References:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-parallelization

NEW QUESTION 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 15
- (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 18
- (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 21
- (Exam Topic 3)
The data engineering team manages Azure HDInsight clusters. The team spends a large amount of time creating and destroying clusters daily because most of the data pipeline process runs in minutes.
You need to implement a solution that deploys multiple HDInsight clusters with minimal effort. What should you implement?

A. Azure Databricks
B. Azure Traffic Manager
C. Azure Resource Manager templates
D. Ambari web user interface

Answer: C

Explanation:
A Resource Manager template makes it easy to create the following resources for your application in a single, coordinated operation:
● HDInsight clusters and their dependent resources (such as the default storage account).
● Other resources (such as Azure SQL Database to use Apache Sqoop).
In the template, you define the resources that are needed for the application. You also specify deployment parameters to input values for different environments.
The template consists of JSON and expressions that you use to construct values for your deployment.
References:
https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hadoop-create-linux-clusters-arm-templates

NEW QUESTION 22
- (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.

Components		Answer Area		
Control Database		•	Task	Component
Azure Service Bus		•	Execution results and diagnostics	Component
Azure Storage		•	Job launcher and tracker	Component
Job Service		•	Job metadata and state	Component

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 23
- (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 27
- (Exam Topic 3)
You are creating a managed data warehouse solution on Microsoft Azure.
You must use PolyBase to retrieve data from Azure Blob storage that resides in parquet format and toad the data into a large table called FactSalesOrderDetails.
You need to configure Azure SQL Data Warehouse to receive the data.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Create an external file format to map the parquet files.	
Load the data to a staging table.	
Create the external table FactSalesOrderDetails.	
Enable Transparent Data Encryption.	
Create an external data source for Azure Blob storage.	⟩
Create a master key on database	⟨
Configure PolyBase to use Azure Blob storage.	

A. Mastered
B. Not Mastered

Answer: A

Explanation:
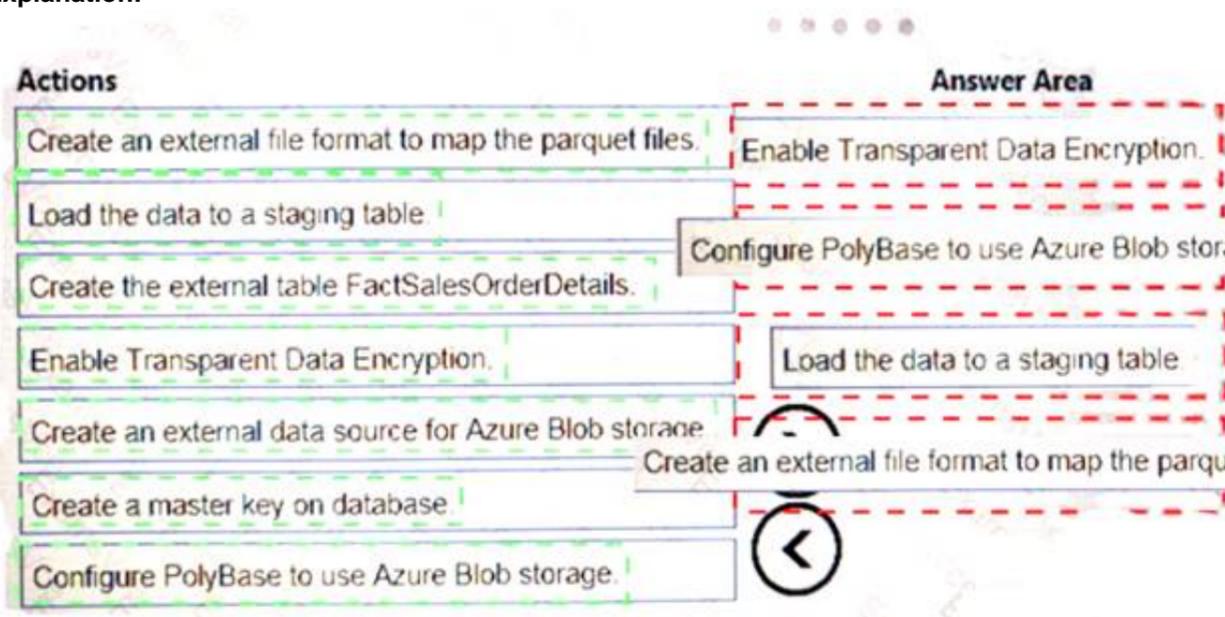
NEW QUESTION 31
- (Exam Topic 3)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.
You develop a data ingestion process that will import data to a Microsoft 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 36
- (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.

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 38
- (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.
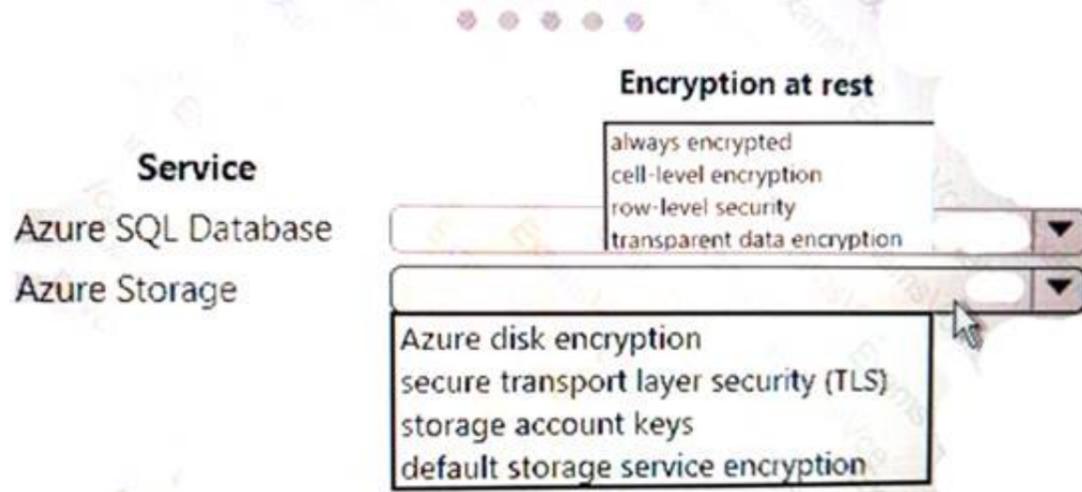
A. Mastered
B. Not Mastered

Answer: A

Explanation:
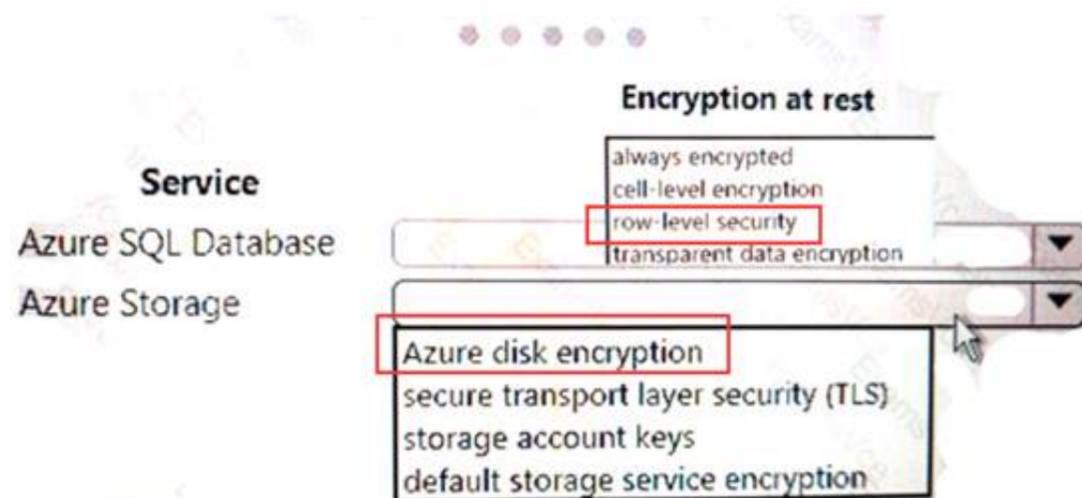
NEW QUESTION 40
......

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