Exam Questions DP-201

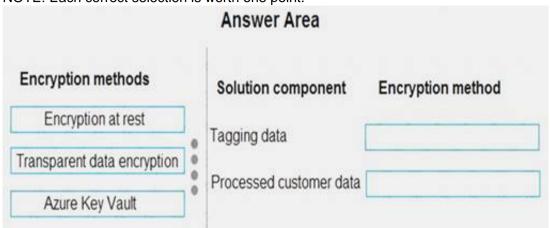Designing an Azure Data Solution

https://www.2passeasy.com/dumps/DP-201/

NEW QUESTION 1
- (Exam Topic 2)
You need to design the encryption strategy for the tagging data and customer data.
What should you recommend? To answer, drag the appropriate setting to the correct drop targets. Each source 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:
All cloud data must be encrypted at rest and in transit. Box 1: Transparent data encryption
Encryption of the database file is performed at the page level. The pages in an encrypted database are encrypted before they are written to disk and decrypted when read into memory.
Box 2: Encryption at rest
Encryption at Rest is the encoding (encryption) of data when it is persisted. References:
https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption?view= https://docs.microsoft.com/en-us/azure/security/azure-security-encryption-atrest

NEW QUESTION 2
- (Exam Topic 2)
You need to design a backup solution for the processed customer data. What should you include in the design?

A. AzCopy
B. AdlCopy
C. Geo-Redundancy
D. Geo-Replication

Answer: C

Explanation:
Scenario: All data must be backed up in case disaster recovery is required.
Geo-redundant storage (GRS) is designed to provide at least 99.99999999999999% (16 9's) durability of objects over a given year by replicating your data to a secondary region that is hundreds of miles away from
the primary region. If your storage account has GRS enabled, then your data is durable even in the case of a complete regional outage or a disaster in which the primary region isn't recoverable. References:
https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy-grs

NEW QUESTION 3
- (Exam Topic 3)
You need to design the disaster recovery solution for customer sales data analytics.
Which three actions should you recommend? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Provision multiple Azure Databricks workspaces in separate Azure regions.
B. Migrate users, notebooks, and cluster configurations from one workspace to another in the same region.
C. Use zone redundant storage.
D. Migrate users, notebooks, and cluster configurations from one region to another.
E. Use Geo-redundant storage.
F. Provision a second Azure Databricks workspace in the same region.

Answer: ADE

Explanation:
Scenario: The analytics solution for customer sales data must be available during a regional outage. To create your own regional disaster recovery topology for databricks, follow these requirements:
1. Provision multiple Azure Databricks workspaces in separate Azure regions
2. Use Geo-redundant storage.
3. Once the secondary region is created, you must migrate the users, user folders, notebooks, cluster configuration, jobs configuration, libraries, storage, init scripts, and reconfigure access control.
Note: Geo-redundant storage (GRS) is designed to provide at least 99.99999999999999% (16 9's) durability of objects over a given year by replicating your data to a secondary region that is hundreds of miles away from the primary region. If your storage account has GRS enabled, then your data is durable even in the case of a complete regional outage or a disaster in which the primary region isn't recoverable.

References:
https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy-grs

NEW QUESTION 4
- (Exam Topic 3)
You need to recommend an Azure SQL Database service tier. What should you recommend?

A. Business Critical
B. General Purpose
C. Premium
D. Standard
E. Basic

Answer: C

Explanation:
The data engineers must set the SQL Data Warehouse compute resources to consume 300 DWUs. Note: There are three architectural models that are used in Azure SQL Database:
- General Purpose/Standard
- Business Critical/Premium
- Hyperscale

NEW QUESTION 5
- (Exam Topic 4)
A company purchases IoT devices to monitor manufacturing machinery. The company uses an IoT appliance to communicate with the IoT devices.
The company must be able to monitor the devices in real-time. You need to design the solution.
What should you recommend?

A. Azure Stream Analytics cloud job using Azure PowerShell
B. Azure Analysis Services using Azure Portal
C. Azure Data Factory instance using Azure Portal
D. Azure Analysis Services using Azure PowerShell

Answer: D

NEW QUESTION 6
- (Exam Topic 4)
You design data engineering solutions for a company.
You must integrate on-premises SQL Server data into an Azure solution that performs Extract-Transform-Load (ETL) operations have the following requirements:
- Develop a pipeline that can integrate data and run notebooks.
- Develop notebooks to transform the data.
- Load the data into a massively parallel processing database for later analysis. You need to recommend a solution.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Requirement	Service
Integrate the on-premises data into the cloud.	▼ Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch
Develop notebooks to transform the data.	▼ Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch
Run notebooks.	▼ Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch
Load the data.	▼ Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch
Store the transformed data.	▼ Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Requirement	Service
Integrate the on-premises data into the cloud.	▼ Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch
Develop notebooks to transform the data.	▼ Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch
Run notebooks.	▼ Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch
Load the data.	▼ Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch
Store the transformed data.	▼ Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch

NEW QUESTION 7
- (Exam Topic 4)
You need to design the unauthorized data usage detection system. What Azure service should you include in the design?

A. Azure Databricks
B. Azure SQL Data Warehouse
C. Azure Analysis Services
D. Azure Data Factory

Answer: B

NEW QUESTION 8
- (Exam Topic 4)
You plan to migrate data to Azure SQL Database.
The database must remain synchronized with updates to Microsoft Azure and SQL Server. You need to set up the database as a subscriber.
What should you recommend?

A. Azure Data Factory
B. SQL Server Data Tools
C. Data Migration Assistant
D. SQL Server Agent for SQL Server 2017 or later
E. SQL Server Management Studio 17.9.1 or later

Answer: E

Explanation:
To set up the database as a subscriber we need to configure database replication. You can use SQL Server Management Studio to configure replication. Use the latest versions of SQL Server Management Studio in order to be able to use all the features of Azure SQL Database.
References:
https://www.sqlshack.com/sql-server-database-migration-to-azure-sql-database-using-sql-server-transactionalrep
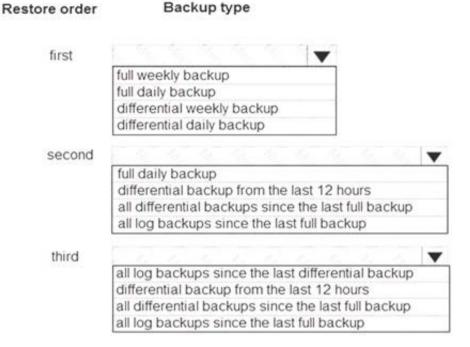
NEW QUESTION 9
- (Exam Topic 4)
You are designing a recovery strategy for your Azure SQL Databases.
The recovery strategy must use default automated backup settings. The solution must include a Point-in time restore recovery strategy.
You need to recommend which backups to use and the order in which to restore backups.

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

Restore order	Backup type

first

| full weekly backup |
| full daily backup |
| differential weekly backup |
| differential daily backup |

second

| full daily backup |
| differential backup from the last 12 hours |
| all differential backups since the last full backup |
| all log backups since the last full backup |

third

| all log backups since the last differential backup |
| differential backup from the last 12 hours |
| all differential backups since the last full backup |
| all log backups since the last full backup |

A. Mastered
B. Not Mastered

Answer: A

Explanation:
All Basic, Standard, and Premium databases are protected by automatic backups. Full backups are taken every week, differential backups every day, and log backups every 5 minutes.
References:
https://azure.microsoft.com/sv-se/blog/azure-sql-database-point-in-time-restore/

NEW QUESTION 10
- (Exam Topic 4)
You are designing a solution for a company. The solution will use model training for objective classification. You need to design the solution.
What should you recommend?

A. an Azure Cognitive Services application
B. a Spark Streaming job
C. interactive Spark queries
D. Power BI models
E. a Spark application that uses Spark MLib.

Answer: E

Explanation:
Spark in SQL Server big data cluster enables AI and machine learning.
You can use Apache Spark MLlib to create a machine learning application to do simple predictive analysis on an open dataset.
MLib is a core Spark library that provides many utilities useful for machine learning tasks, including utilities that are suitable for:
🔵 Classification
🔵 Regression
🔵 Clustering
🔵 Topic modeling
🔵 Singular value decomposition (SVD) and principal component analysis (PCA)
🔵 Hypothesis testing and calculating sample statistics
References:
https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-machine-learning-mllib-ipython

NEW QUESTION 10
- (Exam Topic 4)
You are designing a data processing solution that will run as a Spark job on an HDInsight cluster. The solution will be used to provide near real-time information about online ordering for a retailer.
The solution must include a page on the company intranet that displays summary information. The summary information page must meet the following requirements:
🔵 Display a summary of sales to date grouped by product categories, price range, and review scope.
🔵 Display sales summary information including total sales, sales as compared to one day ago and sales as compared to one year ago.
🔵 Reflect information for new orders as quickly as possible. You need to recommend a design for the solution.
What should you recommend? To answer, select the appropriate configuration in the answer area.

Use case	Technology
Data abstraction	▼ Resilient Distributed Dataset (RDD) Dataset DataFrame
Data format	▼ Avro parquet

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: DataFrame
DataFrames
Best choice in most situations.
Provides query optimization through Catalyst. Whole-stage code generation.
Direct memory access.
Low garbage collection (GC) overhead.
Not as developer-friendly as DataSets, as there are no compile-time checks or domain object programming. Box 2: parquet
The best format for performance is parquet with snappy compression, which is the default in Spark 2.x. Parquet stores data in columnar format, and is highly optimized in Spark.

NEW QUESTION 11
- (Exam Topic 4)
You are evaluating data storage solutions to support a new application.
You need to recommend a data storage solution that represents data by using nodes and relationships in graph structures.
Which data storage solution should you recommend?

A. Blob Storage
B. Cosmos DB
C. Data Lake Store
D. HDInsight

Answer: B

Explanation:
For large graphs with lots of entities and relationships, you can perform very complex analyses very quickly. Many graph databases provide a query language that you can use to traverse a network of relationships efficiently.
Relevant Azure service: Cosmos DB
References:
https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview

NEW QUESTION 16
- (Exam Topic 4)
You have a Windows-based solution that analyzes scientific data. You are designing a cloud-based solution that performs real-time analysis of the data.
You need to design the logical flow for the solution.
Which two actions should you recommend? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Send data from the application to an Azure Stream Analytics job.
B. Use an Azure Stream Analytics job on an edge devic
C. Ingress data from an Azure Data Factory instance and build queries that output to Power BI.
D. Use an Azure Stream Analytics job in the clou
E. Ingress data from the Azure Event Hub instance and build queries that output to Power BI.
F. Use an Azure Stream Analytics job in the clou
G. Ingress data from an Azure Event Hub instance and build queries that output to Azure Data Lake Storage.
H. Send data from the application to Azure Data Lake Storage.
I. Send data from the application to an Azure Event Hub instance.

Answer: CF

Explanation:
Stream Analytics has first-class integration with Azure data streams as inputs from three kinds of resources: Azure Event Hubs
Azure IoT Hub Azure Blob storage References:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-define-inputs

NEW QUESTION 20
- (Exam Topic 4)
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.

You are designing an Azure SQL Database that will use elastic pools. You plan to store data about customers in a table. Each record uses a value for CustomerID. You need to recommend a strategy to partition data based on values in CustomerID. Proposed Solution: Separate data into customer regions by using vertical partitioning. Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Vertical partitioning is used for cross-database queries. Instead we should use Horizontal Partitioning, which also is called charding.
References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-query-overview

NEW QUESTION 21
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DP-201 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DP-201 Product From:

https://www.2passeasy.com/dumps/DP-201/

Money Back Guarantee

DP-201 Practice Exam Features:

* DP-201 Questions and Answers Updated Frequently

* DP-201 Practice Questions Verified by Expert Senior Certified Staff

* DP-201 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* DP-201 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year