[May-2020-NewFree Microsoft 224Q DP-200 PDF and VCE Dumps Braindump2go Offer[181-192

 May/2020 New Braindump2go DP-200 Exam Dumps with PDF and VCE Free Updated Today! Following are some new DP-200 Exam Questions,QUESTION 181You 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: ADEExplanation: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 regions2. 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-grsQUESTION 182Note: 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 HDInsight/Hadoop cluster solution that uses Azure Data Lake Gen1 Storage.The solution requires POSIX permissions and enables diagnostics logging for auditing.You need to recommend solutions that optimize storage.Proposed Solution: Ensure that files stored are larger than 250MB.Does the solution meet the goal?A.   YesB.   NoAnswer: AExplanation: Depending on what services and workloads are using the data, a good size to consider for files is 256 MB or greater. If the file sizes cannot be batched when landing in Data Lake Storage Gen1, you can have a separate compaction job that combines these files into larger ones.Note: POSIX permissions and auditing in Data Lake Storage Gen1 comes with an overhead that becomes apparent when working with numerous small files. As a best practice, you must batch your data into larger files versus writing thousands or millions of small files to Data Lake Storage Gen1. Avoiding small file sizes can have multiple benefits, such as:Lowering the authentication checks across multiple files Reduced open file connectionsFaster copying/replicationFewer files to process when updating Data Lake Storage Gen1 POSIX permissions References:

https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-best-practicesQUESTION 183Note: 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 HDInsight/Hadoop cluster solution that uses Azure Data Lake Gen1 Storage.The solution requires POSIX permissions and enables diagnostics logging for auditing.You need to recommend solutions that optimize storage.Proposed Solution: Implement compaction jobs to combine small files into larger files.Does the solution meet the goal?A.   YesB. NoAnswer: AExplanation:Depending on what services and workloads are using the data, a good size to consider for files is 256 MB or greater. If the file sizes cannot be batched when landing in Data Lake Storage Gen1, you can have a separate compaction job that combines these files into larger ones.Note: POSIX permissions and auditing in Data Lake Storage Gen1 comes with an overhead that becomes apparent when working with numerous small files. As a best practice, you must batch your data into larger files versus writing thousands or millions of small files to Data Lake Storage Gen1. Avoiding small file sizes can have multiple benefits, such as:Lowering the authentication checks across multiple files Reduced open file connectionsFaster copying/replicationFewer files to process when updating Data Lake Storage Gen1 POSIX permissionsReferences:

https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-best-practicesQUESTION 184Note: 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 HDInsight/Hadoop cluster solution that uses Azure Data Lake Gen1 Storage. The solution requires

POSIX permissions and enables diagnostics logging for auditing.You need to recommend solutions that optimize storage.Proposed Solution: Ensure that files stored are smaller than 250MB.Does the solution meet the goal?A.   YesB.   NoAnswer: BExplanation: Ensure that files stored are larger, not smaller than 250MB.You can have a separate compaction job that combines these files into larger ones.Note: The file POSIX permissions and auditing in Data Lake Storage Gen1 comes with an overhead that becomes apparent when working with numerous small files. As a best practice, you must batch your data into larger files versus writing thousands or millions of small files to Data Lake Storage Gen1. Avoiding small file sizes can have multiple benefits, such as: Lowering the authentication checks across multiple files Reduced open file connectionsFaster copying/replicationFewer files to process when updating Data Lake Storage Gen1 POSIX permissionsReferences:

https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-best-practicesQUESTION 185Note: 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.   YesB.   NoAnswer: BExplanation: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

QUESTION 186Note: 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 horizontal partitioning.Does the solution meet the goal?A.   YesB.   NoAnswer: BExplanation:We should use Horizontal Partitioning through Sharding, not divide through regions.Note: Horizontal Partitioning - Sharding: Data is partitioned horizontally to distribute rows across a scaled out data tier. With this approach, the schema is identical on all participating databases. This approach is also called "sharding". Sharding can be performed and managed using (1) the elastic database tools libraries or (2) self-sharding. An elastic query is used to query or compile reports across many shards.References:

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-query-overviewQUESTION 187Note: 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 shards by using horizontal partitioning.Does the solution meet the goal?A.   YesB.   NoAnswer: A Explanation:Horizontal Partitioning - Sharding: Data is partitioned horizontally to distribute rows across a scaled out data tier. With this approach, the schema is identical on all participating databases. This approach is also called "sharding". Sharding can be performed and managed using (1) the elastic database tools libraries or (2) self- sharding. An elastic query is used to query or compile reports across many shards.References:

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-query-overviewQUESTION 188You 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 StorageB.   Cosmos DBC.   Data Lake StoreD.   HDInsightAnswer: BExplanation: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 DBReferences:

https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overviewQUESTION 189You are designing a data processing solution that will implement the lambda architecture pattern. The solution will use Spark running on HDInsight for data processing.You need to recommend a data storage technology for the solution.Which two technologies should you recommend? Each correct answer presents a complete solution.NOTE: Each correct selection is worth one point.A.   Azure Cosmos DBB.   Azure Service BusC.   Azure Storage QueueD.   Apache CassandraE.   Kafka HDInsightAnswer: AEExplanation:

To implement a lambda architecture on Azure, you can combine the following technologies to accelerate real-time big data analytics:Azure Cosmos DB, the industry's first globally distributed, multi-model database service.Apache Spark for Azure HDInsight, a processing framework that runs large-scale data analytics applicationsAzure Cosmos DB change feed, which streams new data to the batch layer for HDInsight to process The Spark to Azure Cosmos DB ConnectorE: You can use Apache Spark to stream data into or out of Apache Kafka on HDInsight using DStreams.References:

https://docs.microsoft.com/en-us/azure/cosmos-db/lambda-architectureQUESTION 190A company manufactures automobile parts. The company installs IoT sensors on manufacturing machinery.You must design a solution that analyzes data from the sensors.You need to recommend a solution that meets the following requirements:- Data must be analyzed in real-time.- Data queries must be deployed using continuous integration.- Data must be visualized by using charts and graphs.- Data must be available for ETL operations in the future.- The solution must support high-volume data ingestion.Which three actions should you recommend? Each correct answer presents part of the solution.NOTE: Each correct selection is worth one point.A.   Use Azure Analysis Services to query the data. Output query results to Power BI.B.   Configure an Azure Event Hub to capture data to Azure Data Lake Storage.C.   Develop an Azure Stream Analytics application that queries the data and outputs to Power BI. Use Azure Data Factory to deploy the Azure Stream Analytics application.D.   Develop an application that sends the IoT data to an Azure Event Hub.E.   Develop an Azure Stream Analytics application that queries the data and outputs to Power BI. Use Azure Pipelines to deploy the Azure Stream Analytics application.F.   Develop an application that sends the IoT data to an Azure Data Lake Storage container.Answer: BCDQUESTION 191You are designing an Azure Databricks interactive cluster.You need to ensure that the cluster meets the following requirements:- Enable auto-termination- Retain cluster configuration indefinitely after cluster termination.What should you recommend?A.   Start the cluster after it is terminated.B.   Pin the clusterC.   Clone the cluster after it is terminated.D.   Terminate the cluster manually at process completion.Answer: BExplanation:To keep an interactive cluster configuration even after it has been terminated for more than 30 days, an administrator can pin a cluster to the cluster list.References:

https://docs.azuredatabricks.net/user-guide/clusters/terminate.htmlQUESTION 192You 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 applicationB.   a Spark Streaming jobC.   interactive Spark queriesD.   Power BI modelsE.   a Spark application that uses Spark MLib.Answer: EExplanation: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.MLlib is a core Spark library that provides many utilities useful for machine learning tasks, including utilities that are suitable for:Classification RegressionClusteringTopic modelingSingular value decomposition (SVD) and principal component analysis (PCA) Hypothesis testing and calculating sample statisticsReferences:

https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-machine-learning-mllib-ipythonResources From:1.2020 Latest Braindump2go DP-200 Exam Dumps (PDF & VCE) Free Share:https://www.braindump2go.com/dp-200.html2.2020 Latest Braindump2go DP-200 PDF and DP-200 VCE Dumps Free Share:

https://drive.google.com/drive/folders/1Lr1phOaAVcbL-_R5O-DFweNoShul0W13?usp=sharing3.2020 Free Braindump2go DP-200 PDF Download:https://www.braindump2go.com/free-online-pdf/DP-200-Dumps(186-196).pdf

https://www.braindump2go.com/free-online-pdf/DP-200-PDF(175-185).pdf

https://www.braindump2go.com/free-online-pdf/DP-200-PDF-Dumps(153-163).pdf

https://www.braindump2go.com/free-online-pdf/DP-200-VCE(142-152).pdf

https://www.braindump2go.com/free-online-pdf/DP-200-VCE-Dumps(164-174).pdfFree Resources from Braindump2go,We Devoted to Helping You 100% Pass All Exams!