[New Exams!Real Exam Questions-Braindump2go DP-201 PDF Dumps 75Q Download

2019/July Braindump2go DP-201 Exam Dumps with PDF and VCE New Updated Today! Following are some new DP-201 Real Exam Questions:1.|2019 Latest Braindump2go DP-201 Exam Dumps (PDF & VCE) Instant Download:https://www.braindump2go.com/dp-201.html2.|2019 Latest Braindump2go DP-201 Exam Questions & Answers Instant Download:https://drive.google.com/drive/folders/1umFAfoENMrqFV_co0v9XQ_IvY1RaVBOm?usp=sharingNew QuestionCase Study 3 BackgroundCurrent environmentThe company has the following virtual machines (VMs):RequirementsStorage and processingYou must be able to use a file system view of data stored in a blob.You must build an architecture that will allow Contoso to use the DB FS filesystem layer over a blob store. The architecture will need to support data files, libraries, and images. Additionally, it must provide a web-based interface to documents that contain runnable command, visualizations, and narrative text such as a notebook.CONT_SQL3 requires an initial scale of 35000 IOPS.CONT_SQL1 and CONT_SQL2 must use the vCore model and should include replicas. The solution must support 8000 IOPS. The storage should be configured to optimized storage for database OLTP workloads.MigrationYou must be able to independently scale compute and storage resources.You must migrate all SQL Server workloads to Azure. You must identify related machines in the on-premises environment, get disk size data usage information.Data from SQL Server must include zone redundant storage.You need to ensure that app components can reside on-premises while interacting with components that run in the Azure public cloud. SAP data must remain on-premises.The Azure Site Recovery (ASR) results should contain per-machine data.Business requirementsYou must design a regional disaster recovery topology.The database backups have regulatory purposes and must be retained for seven years.CONT_SQL1 stores customers sales data that requires ETL operations for data analysis. A solution is required that reads data from SQL, performs ETL, and outputs to Power BI. The solution should use managed clusters to minimize costs. To optimize logistics, Contoso needs to analyze customer sales data to see if certain products are tied to specific times in the year.The analytics solution for customer sales data must be available during a regional outage.Security and auditingContoso requires all corporate computers to enable Windows Firewall. Azure servers should be able to ping other Contoso Azure servers.Employee PII must be encrypted in memory, in motion, and at rest. Any data encrypted by SQL Server must support equality searches, grouping, indexing, and joining on the encrypted data.Keys must be secured by using hardware security modules (HSMs). CONT_SQL3 must not communicate over the default portsCostAll solutions must minimize cost and resources.The organization does not want any unexpected charges.The data engineers must set the SQL Data Warehouse compute resources to consume 300 DWUs.CONT_SQL2 is not fully utilized during non-peak hours. You must minimize resource costs for during non-peak hours.You need to optimize storage for CONT_SQL3. What should you recommend?A. AlwaysOnB. Transactional processingC. GeneralD. Data warehousingCorrect Answer: BExplanation Explanation/Reference:Explanation:CONT_SQL3 with the SQL Server role, 100 GB database size, Hyper-VM to be migrated to Azure VM. The storage should be configured to optimized storage for database OLTP workloads.Azure SQL Database provides three basic in-memory based capabilities (built into the underlying database engine) that can contribute in a meaningful way to performance improvements:In-Memory Online Transactional Processing (OLTP)Clustered columnstore indexes intended primarily for Online Analytical Processing (OLAP) workloads Nonclustered columnstore indexes geared towards Hybrid Transactional/Analytical Processing (HTAP) workloads References:
https://www.databasejournal.com/features/mssql/overview-of-in-memory-technologies-of-azure-sql-database.html New QuestionCase Study 3 BackgroundCurrent environmentThe company has the following virtual machines (VMs):Requirements Storage and processingYou must be able to use a file system view of data stored in a blob.You must build an architecture that will allow Contoso to use the DB FS filesystem layer over a blob store. The architecture will need to support data files, libraries, and images. Additionally, it must provide a web-based interface to documents that contain runnable command, visualizations, and narrative text such as a notebook.CONT_SQL3 requires an initial scale of 35000 IOPS.CONT_SQL1 and CONT_SQL2 must use the vCore model and should include replicas. The solution must support 8000 IOPS. The storage should be configured to optimized storage for database OLTP workloads.MigrationYou must be able to independently scale compute and storage resources.You must migrate all SQL Server workloads to Azure. You must identify related machines in the on-premises environment, get disk size data usage information.Data from SQL Server must include zone redundant storage.You need to ensure that app components can reside on-premises while interacting with components that run in the Azure public cloud. SAP data must remain on-premises.The Azure Site Recovery (ASR) results should contain per-machine data.Business requirementsYou must design a regional disaster recovery topology.The database backups have regulatory purposes and must be retained for seven years.CONT_SQL1 stores customers sales data that requires ETL operations for data analysis. A solution is required that reads data from SQL, performs ETL, and outputs to

Power BI. The solution should use managed clusters to minimize costs. To optimize logistics, Contoso needs to analyze customer sales data to see if certain products are tied to specific times in the year.The analytics solution for customer sales data must be available during a regional outage.Security and auditingContoso requires all corporate computers to enable Windows Firewall. Azure servers should be able to ping other Contoso Azure servers.Employee PII must be encrypted in memory, in motion, and at rest. Any data encrypted by SQL Server must support equality searches, grouping, indexing, and joining on the encrypted data.Keys must be secured by using hardware security modules (HSMs). CONT_SQL3 must not communicate over the default portsCostAll solutions must minimize cost and resources.The organization does not want any unexpected charges.The data engineers must set the SQL Data Warehouse compute resources to consume 300 DWUs.CONT_SQL2 is not fully utilized during non-peak hours. You must minimize resource costs for during non-peak hours.You need to recommend a backup strategy for CONT_SQL1 and CONT_SQL2. What should you recommend?A. Use AzCopy and store the data in Azure.B. Configure Azure SQL Database long-term retention for all databases.C. Configure Accelerated Database Recovery.D. Use DWLoader.Correct Answer: B ExplanationExplanation/Reference:Explanation:Scenario: The database backups have regulatory purposes and must be retained for seven years.New QuestionCase Study 3 BackgroundCurrent environmentThe company has the following virtual machines (VMs): RequirementsStorage and processingYou must be able to use a file system view of data stored in a blob.You must build an architecture that will allow Contoso to use the DB FS filesystem layer over a blob store. The architecture will need to support data files, libraries, and images. Additionally, it must provide a web-based interface to documents that contain runnable command, visualizations, and narrative text such as a notebook.CONT_SQL3 requires an initial scale of 35000 IOPS.CONT_SQL1 and CONT_SQL2 must use the vCore model and should include replicas. The solution must support 8000 IOPS. The storage should be configured to optimized storage for database OLTP workloads.MigrationYou must be able to independently scale compute and storage resources.You must migrate all SQL Server workloads to Azure. You must identify related machines in the on-premises environment, get disk size data usage information.Data from SQL Server must include zone redundant storage.You need to ensure that app components can reside on-premises while interacting with components that run in the Azure public cloud. SAP data must remain on-premises.The Azure Site Recovery (ASR) results should contain per-machine data.Business requirementsYou must design a regional disaster recovery topology.The database backups have regulatory purposes and must be retained for seven years. CONT_SQL1 stores customers sales data that requires ETL operations for data analysis. A solution is required that reads data from SQL, performs ETL, and outputs to Power BI. The solution should use managed clusters to minimize costs. To optimize logistics, Contoso needs to analyze customer sales data to see if certain products are tied to specific times in the year.The analytics solution for customer sales data must be available during a regional outage.Security and auditingContoso requires all corporate computers to enable Windows Firewall. Azure servers should be able to ping other Contoso Azure servers.Employee PII must be encrypted in memory, in motion, and at rest. Any data encrypted by SQL Server must support equality searches, grouping, indexing, and joining on the encrypted data.Keys must be secured by using hardware security modules (HSMs). CONT_SQL3 must not communicate over the default ports CostAll solutions must minimize cost and resources.The organization does not want any unexpected charges.The data engineers must set the SQL Data Warehouse compute resources to consume 300 DWUs.CONT_SQL2 is not fully utilized during non-peak hours. You must minimize resource costs for during non-peak hours.Hotspot QuestionYou need to design network access to the SQL Server data.What should you recommend? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.Hot Area:Correct Answer: ExplanationExplanation/Reference:Explanation:Box 1: 8080 1433 is the default port, but we must change it as CONT_SQL3 must not communicate over the default ports. Because port 1433 is the known standard for SQL Server, some organizations specify that the SQL Server port number should be changed to enhance security.Box 2: SQL Server Configuration ManagerYou can configure an instance of the SQL Server Database Engine to listen on a specific fixed port by using the SQL Server Configuration Manager.References: **https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-a-server-to-listen-on-a-specific-tcp-port?view=sql-server-2017**New QuestionYou 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.Correct Answer: ADEExplanationExplanation/Reference: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 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-grs**New QuestionNote: 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. NoCorrect Answer: A ExplanationExplanation/Reference:Explanation: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-practices**New QuestionNote: 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. NoCorrect Answer: AExplanationExplanation/Reference:Explanation: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 connections Faster 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-practices**New QuestionNote: 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. NoCorrect Answer: B ExplanationExplanation/Reference:Explanation: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 permissions References: **https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-best-practices**New QuestionNote: 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. NoCorrect Answer: BExplanationExplanation/Reference: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 QuestionNote: 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. NoCorrect Answer: BExplanationExplanation/Reference:Explanation: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-overview New QuestionNote: 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. NoCorrect Answer: AExplanationExplanation/Reference: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-overview New QuestionYou 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. HDInsight Correct Answer: BExplanationExplanation/Reference: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!!!RECOMMEND!!!1.|2019 Latest Braindump2go DP-201 Exam Dumps (PDF & VCE) Instant Download:https://www.braindump2go.com/dp-201.html2.|2019 Latest Braindump2go DP-201 Study Guide Video Instant Download: YouTube Video: YouTube.com/watch?v=8h9yuqa-Vb8