

## [New Exams!Braindump2go DP-201 Exam PDF and VCE Dumps 75Q Free Offer

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.html>2.|2019 Latest Braindump2go DP-201 Exam Questions & Answers Instant Download:[https://drive.google.com/drive/folders/1umFAfoENMrqFV\\_co0v9XQ\\_IvY1RaVBOM?usp=sharing](https://drive.google.com/drive/folders/1umFAfoENMrqFV_co0v9XQ_IvY1RaVBOM?usp=sharing)New QuestionCase

Study 2 Requirements BusinessThe company identifies the following business requirements:You must transfer all images and customer data to cloud storage and remove on-premises servers. You must develop an analytical processing solution for transforming customer data.You must develop an image object and color tagging solution. Capital expenditures must be minimized. Cloud resource costs must be minimized.TechnicalThe solution has the following technical requirements:Tagging data must be uploaded to the cloud from the New York office location.Tagging data must be replicated to regions that are geographically close to company office locations. Image data must be stored in a single data store at minimum cost.Customer data must be analyzed using managed Spark clusters. Power BI must be used to visualize transformed customer data. All data must be backed up in case disaster recovery is required.Security and optimizationAll cloud data must be encrypted at rest and in transit. The solution must support: parallel processing of customer datahyper-scale storage of imagesglobal region data replication of processed image dataYou need to recommend a solution for storing the image tagging data. What should you recommend?A. Azure File StorageB. Azure Cosmos DBC. Azure Blob StorageD. Azure SQL DatabaseE. Azure SQL Data WarehouseCorrect Answer: CExplanation

Explanation/Reference:Explanation:Image data must be stored in a single data store at minimum cost.Note: Azure Blob storage is Microsoft's object storage solution for the cloud. Blob storage is optimized for storing massive amounts of unstructured data. Unstructured data is data that does not adhere to a particular data model or definition, such as text or binary data.Blob storage is designed for:Serving images or documents directly to a browser. Storing files for distributed access.Streaming video and audio. Writing to log files.Storing data for backup and restore, disaster recovery, and archiving. Storing data for analysis by an on-premises or Azure-hosted service. References:<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blobs-introduction>New

QuestionCase Study 2 Requirements BusinessThe company identifies the following business requirements:You must transfer all images and customer data to cloud storage and remove on-premises servers. You must develop an analytical processing solution for transforming customer data.You must develop an image object and color tagging solution. Capital expenditures must be minimized. Cloud resource costs must be minimized.TechnicalThe solution has the following technical requirements:Tagging data must be uploaded to the cloud from the New York office location.Tagging data must be replicated to regions that are geographically close to company office locations. Image data must be stored in a single data store at minimum cost.Customer data must be analyzed using managed Spark clusters. Power BI must be used to visualize transformed customer data. All data must be backed up in case disaster recovery is required.Security and optimizationAll cloud data must be encrypted at rest and in transit. The solution must support: parallel processing of customer datahyper-scale storage of imagesglobal region data replication of processed image data You need to design the solution for analyzing customer data. What should you recommend?A. Azure DatabricksB. Azure Data Lake Storage C. Azure SQL Data WarehouseD. Azure Cognitive ServicesE. Azure BatchCorrect Answer: AExplanation

Explanation/Reference:Explanation:Customer data must be analyzed using managed Spark clusters. You create spark clusters through Azure Databricks.References:

<https://docs.microsoft.com/en-us/azure/azure-databricks/quickstart-create-databricks-workspace-portal>New QuestionCase Study 2 Requirements BusinessThe company identifies the following business requirements:You must transfer all images and customer data to cloud storage and remove on-premises servers. You must develop an analytical processing solution for transforming customer data.You must develop an image object and color tagging solution. Capital expenditures must be minimized. Cloud resource costs must be minimized.TechnicalThe solution has the following technical requirements:Tagging data must be uploaded to the cloud from the New York office location.Tagging data must be replicated to regions that are geographically close to company office locations. Image data must be stored in a single data store at minimum cost.Customer data must be analyzed using managed Spark clusters. Power BI must be used to visualize transformed customer data. All data must be backed up in case disaster recovery is required.Security and optimizationAll cloud data must be encrypted at rest and in transit. The solution must support: parallel processing of customer datahyper-scale storage of imagesglobal region data replication of processed image data You need to recommend a solution for storing customer data. What should you recommend?A. Azure SQL Data WarehouseB. Azure Stream AnalyticsC. Azure DatabricksD. Azure SQL DatabaseCorrect Answer: CExplanationExplanation/Reference:Explanation:From the scenario:Customer data must be analyzed using managed Spark clusters.All cloud data must be encrypted at rest and in transit. The solution must support: parallel processing of customer data. References:

<https://www.microsoft.com/developerblog/2019/01/18/running-parallel-apache-spark-notebook-workloads-on-azure-databricks/>

**Case Study 2 Requirements Business**The company identifies the following business requirements: You must transfer all images and customer data to cloud storage and remove on-premises servers. You must develop an analytical processing solution for transforming customer data. You must develop an image object and color tagging solution. Capital expenditures must be minimized. Cloud resource costs must be minimized. **Technical**The solution has the following technical requirements: Tagging data must be uploaded to the cloud from the New York office location. Tagging data must be replicated to regions that are geographically close to company office locations. Image data must be stored in a single data store at minimum cost. Customer data must be analyzed using managed Spark clusters. Power BI must be used to visualize transformed customer data. All data must be backed up in case disaster recovery is required. **Security and optimization**All cloud data must be encrypted at rest and in transit. The solution must support: parallel processing of customer data hyper-scale storage of images global region data replication of processed image data 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 **Correct Answer: C** **Explanation** Explanation/Reference: **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>

**Case Study 2 Requirements Business**The company identifies the following business requirements: You must transfer all images and customer data to cloud storage and remove on-premises servers. You must develop an analytical processing solution for transforming customer data. You must develop an image object and color tagging solution. Capital expenditures must be minimized. Cloud resource costs must be minimized. **Technical**The solution has the following technical requirements: Tagging data must be uploaded to the cloud from the New York office location. Tagging data must be replicated to regions that are geographically close to company office locations. Image data must be stored in a single data store at minimum cost. Customer data must be analyzed using managed Spark clusters. Power BI must be used to visualize transformed customer data. All data must be backed up in case disaster recovery is required. **Security and optimization**All cloud data must be encrypted at rest and in transit. The solution must support: parallel processing of customer data hyper-scale storage of images global region data replication of processed image data You plan to use an Azure SQL data warehouse to store the customer data. You need to recommend a disaster recovery solution for the data warehouse. What should you include in the recommendation? A. AzCopy B. Read-only replicas C. AdlCopy D. Geo-Redundant backups **Correct Answer: D** **Explanation** Explanation/Reference: **Explanation:**

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/backup-and-restore>

**Case Study 2 Requirements Business**The company identifies the following business requirements: You must transfer all images and customer data to cloud storage and remove on-premises servers. You must develop an analytical processing solution for transforming customer data. You must develop an image object and color tagging solution. Capital expenditures must be minimized. Cloud resource costs must be minimized. **Technical**The solution has the following technical requirements: Tagging data must be uploaded to the cloud from the New York office location. Tagging data must be replicated to regions that are geographically close to company office locations. Image data must be stored in a single data store at minimum cost. Customer data must be analyzed using managed Spark clusters. Power BI must be used to visualize transformed customer data. All data must be backed up in case disaster recovery is required. **Security and optimization**All cloud data must be encrypted at rest and in transit. The solution must support: parallel processing of customer data hyper-scale storage of images global region data replication of processed image data **Drag and Drop Question**You need to design the image processing solution to meet the optimization requirements for image tag data. What should you configure? 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. **Select and Place:** **Correct Answer:** **Explanation** Explanation/Reference: **Explanation:** Tagging data must be uploaded to the cloud from the New York office location. Tagging data must be replicated to regions that are geographically close to company office locations. **New Question** **Case Study 2 Requirements Business**The company identifies the following business requirements: You must transfer all images and customer data to cloud storage and remove on-premises servers. You must develop an analytical processing solution for transforming customer data. You must develop an image object and color tagging solution. Capital expenditures must be minimized. Cloud resource costs must be minimized. **Technical**The solution has the following technical requirements: Tagging data must be uploaded to the cloud from the New York office location. Tagging data must be replicated to regions that are geographically close to company office locations. Image data must be stored in a single data store at minimum cost. Customer data must be analyzed using

managed Spark clusters. Power BI must be used to visualize transformed customer data. All data must be backed up in case disaster recovery is required. Security and optimization All cloud data must be encrypted at rest and in transit. The solution must support: parallel processing of customer data hyper-scale storage of images global region data replication of processed image data Hotspot Question You need to design the image processing and storage solutions. What should you recommend? To answer, select the appropriate configuration in the answer area. NOTE: Each correct selection is worth one point. Hot Area: Correct Answer: Explanation Explanation/Reference: Explanation: From the scenario: The company identifies the following business requirements: You must transfer all images and customer data to cloud storage and remove on-premises servers. You must develop an image object and color tagging solution. The solution has the following technical requirements: Image data must be stored in a single data store at minimum cost. All data must be backed up in case disaster recovery is required. All cloud data must be encrypted at rest and in transit. The solution must support: hyper-scale storage of images References:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/batch-processing>

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tier-hyperscale> New Question Case Study 2

Requirements Business The company identifies the following business requirements: You must transfer all images and customer data to cloud storage and remove on-premises servers. You must develop an analytical processing solution for transforming customer data. You must develop an image object and color tagging solution. Capital expenditures must be minimized. Cloud resource costs must be minimized. Technical The solution has the following technical requirements: Tagging data must be uploaded to the cloud from the New York office location. Tagging data must be replicated to regions that are geographically close to company office locations. Image data must be stored in a single data store at minimum cost. Customer data must be analyzed using managed Spark clusters. Power BI must be used to visualize transformed customer data. All data must be backed up in case disaster recovery is required. Security and optimization All cloud data must be encrypted at rest and in transit. The solution must support: parallel processing of customer data hyper-scale storage of images global region data replication of processed image data Drag and Drop Question 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. Select and Place: Correct Answer: Explanation Explanation/Reference: 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=sql-server-](https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption?view=sql-server-2017)

[2017 https://docs.microsoft.com/en-us/azure/security/azure-security-encryption-atrest](https://docs.microsoft.com/en-us/azure/security/azure-security-encryption-atrest) New Question Case Study 3 Background Current environment The company has the following virtual machines (VMs): Requirements Storage and processing You 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. Migration You 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 requirements You 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 auditing Contoso 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 Cost All 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 design a solution to meet the SQL Server storage requirements for CONT\_SQL3. Which type of disk should you recommend? A. Standard SSD Managed Disk B. Premium SSD Managed Disk C. Ultra SSD Managed Disk Correct Answer: C Explanation/Reference: Explanation: CONT\_SQL3 requires an initial scale of 35000 IOPS. Ultra SSD Managed Disk Offerings The following table provides a comparison of ultra solid-state-drives (SSD) (preview), premium SSD, standard SSD, and standard hard disk drives (HDD) for managed disks to help you decide what to use. References: <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disks-types>

New Question Case Study 3 Background Current environment The company has the following virtual machines (VMs): Requirements Storage and processing You 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. Migration You 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 requirements You 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 auditing Contoso 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 Cost All 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 an Azure SQL Database service tier. What should you recommend? A. Business Critical B. General Purpose C. Premium D. Standard E. Basic Correct Answer: C Explanation/Reference: 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 Incorrect Answers: A: Business Critical service tier is designed for the applications that require low-latency responses from the underlying SSD storage (1-2 ms in average), fast recovery if the underlying infrastructure fails, or need to off-load reports, analytics, and read-only queries to the free of charge readable secondary replica of the primary database. References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tier-business-critical>

New Question Case Study 3 Background Current environment The company has the following virtual machines (VMs): Requirements Storage and processing You 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. Migration You 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 requirements You 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 auditing Contoso 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. Cost All 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 the appropriate storage and processing solution? What should you recommend?

A. Enable auto-shrink on the database.  
B. Flush the blob cache using Windows PowerShell.  
C. Enable Apache Spark RDD (RDD) caching.  
D. Enable Databricks IO (DBIO) caching.  
E. Configure the reading speed using Azure Data Studio.

Correct Answer: C

Explanation/Reference: Explanation: Scenario: You 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. Databricks File System (DBFS) is a distributed file system installed on Azure Databricks clusters. Files in DBFS persist to Azure Blob storage, so you won't lose data even after you terminate a cluster. The Databricks Delta cache, previously named Databricks IO (DBIO) caching, accelerates data reads by creating copies of remote files in nodes' local storage using a fast intermediate data format. The data is cached automatically whenever a file has to be fetched from a remote location. Successive reads of the same data are then performed locally, which results in significantly improved reading speed.

References:

<https://docs.databricks.com/delta/delta-cache.html#delta-cache!!!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](https://www.youtube.com/watch?v=8h9yuqa-Vb8)