

[New Exams!Free DP-200 PDF Dumps and DP-200 VCE Offered by Braindump2go[Q1-Q11

2019/July Braindump2go DP-200 Exam Dumps with PDF and VCE New Updated Today! Following are some new DP-200 Real Exam Questions:New QuestionYou 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 secondYou need to create an HDInsight cluster to meet the requirements.Which type of cluster should you create?A. Interactive QueryB. Apache HadoopC. Apache HBaseD. Apache SparkAnswer: DExplanation:Lambda Architecture with Azure:Azure offers you a combination of following technologies to accelerate real-time big data analytics:1. Azure Cosmos DB, a globally distributed and multi-model database service.2. Apache Spark for Azure HDInsight, a processing framework that runs large-scale data analytics applications.3. Azure Cosmos DB change feed, which streams new data to the batch layer for HDInsight to process.4. 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-db/>New QuestionYou develop data engineering solutions for a company. The company has on-premises Microsoft SQL Server databases at multiple locations.The company must integrate data with Microsoft Power BI and Microsoft Azure Logic Apps. The solution must avoid single points of failure during connection and transfer to the cloud. The solution must also minimize latency.You need to secure the transfer of data between on-premises databases and Microsoft Azure.What should you do?A. Install a standalone on-premises Azure data gateway at each locationB. Install an on-premises data gateway in personal mode at each locationC. Install an Azure on-premises data gateway at the primary locationD. Install an Azure on-premises data gateway as a cluster at each locationAnswer: DExplanation:You can create high availability clusters of On-premises data gateway installations, to ensure your organization can access on-premises data resources used in Power BI reports and dashboards. Such clusters allow gateway administrators to group gateways to avoid single points of failure in accessing on-premises data resources. The Power BI service always uses the primary gateway in the cluster, unless it's not available. In that case, the service switches to the next gateway in the cluster, and so on.References:<https://docs.microsoft.com/en-us/power-bi/service-gateway-high-availability-clusters>New QuestionYou are a data architect. The data engineering team needs to configure a synchronization of data between an on-premises Microsoft SQL Server database to Azure SQL Database.Ad-hoc and reporting queries are being overutilized the on-premises production instance. The synchronization process must:- Perform an initial data synchronization to Azure SQL Database with minimal downtime- Perform bi-directional data synchronization after initial synchronizationYou need to implement this synchronization solution.Which synchronization method should you use?A. transactional replicationB. Data Migration Assistant (DMA)C. backup and restoreD. SQL Server Agent jobE. Azure SQL Data SyncAnswer: EExplanation:SQL Data Sync is a service built on Azure SQL Database that lets you synchronize the data you select bi-directionally across multiple SQL databases and SQL Server instances.With Data Sync, you can keep data synchronized between your on-premises databases and Azure SQL databases to enable hybrid applications.Compare Data Sync with Transactional Replication References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-sync-data>New QuestionAn application will use Microsoft Azure Cosmos DB as its data solution. The application will use the Cassandra API to support a column-based database type that uses containers to store items.You need to provision Azure Cosmos DB. Which container name and item name should you use? Each correct answer presents part of the solutions.NOTE: Each correct answer selection is worth one point.A. collectionB. rowsC. graphD. entitiesE. tableAnswer: BExplanation:B: Depending on the choice of the API, an Azure Cosmos item can represent either a document in a collection, a row in a table or a node/edge in a graph. The following table shows the mapping between API-specific entities to an Azure Cosmos item: E: An Azure Cosmos container is specialized into API-specific entities as follows: References:<https://docs.microsoft.com/en-us/azure/cosmos-db/databases-containers-items>New QuestionA company has a SaaS solution that uses Azure SQL Database with elastic pools. The solution contains a dedicated database for each customer organization.Customer organizations have peak usage at different periods during the year.You need to implement the Azure SQL Database elastic pool to minimize cost.Which option or options should you configure?A. Number of transactions onlyB. eDTUs per database onlyC. Number of databases onlyD. CPU usage onlyE. eDTUs and max data sizeAnswer: EExplanation:The best size for a pool depends on the aggregate resources needed for all databases in the pool. This involves determining the following: Maximum resources utilized by all databases in the pool (either maximum DTUs or maximum vCores depending on your choice of

resourcing model).Maximum storage bytes utilized by all databases in the pool.Note: Elastic pools enable the developer to purchase resources for a pool shared by multiple databases to accommodate unpredictable periods of usage by individual databases. You can configure resources for the pool based either on the DTU-based purchasing model or the vCore-based purchasing model.References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-pool>New QuestionA 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 databaseB. Azure SQL Data WarehouseC. Azure Cosmos DBD. Azure SQL Database Managed InstanceAnswer: DExplanation: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 QuestionThe 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 DatabricksB. Azure Traffic ManagerC. Azure Resource Manager templatesD. Ambari web user interfaceAnswer: CExplanation: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 QuestionYou are the data engineer for your company. An application uses a NoSQL database to store data. The database uses the key-value and wide-column NoSQL database type.Developers need to access data in the database using an API.You need to determine which API to use for the database model and type.Which two APIs should you use? Each correct answer presents a complete solution.NOTE: Each correct selection is worth one point.A. Table APIB. MongoDB APIC. Gremlin APID. SQL APIE. Cassandra APIAnswer: BEEExplanation:B: Azure Cosmos DB is the globally distributed, multimodel database service from Microsoft for mission-critical applications. It is a multimodel database and supports document, key-value, graph, and columnar data models.E: Wide-column stores store data together as columns instead of rows and are optimized for queries over large datasets. The most popular are Cassandra and HBase.References:**<https://docs.microsoft.com/en-us/azure/cosmos-db/graph-introduction>**

<https://www.mongodb.com/scale/types-of-nosql-databases>New QuestionA company is designing a hybrid solution to synchronize data and on-premises Microsoft SQL Server database to Azure SQL Database.You must perform an assessment of databases to determine whether data will move without compatibility issues. You need to perform the assessment.Which tool should you use?A. SQL Server Migration Assistant (SSMA)B. Microsoft Assessment and Planning ToolkitC. SQL Vulnerability Assessment (VA)D. Azure SQL Data SyncE. Data Migration Assistant (DMA)Answer: EExplanation:The Data Migration Assistant (DMA) helps you upgrade to a modern data platform by detecting compatibility issues that can impact database functionality in your new version of SQL Server or Azure SQL Database. DMA recommends performance and reliability improvements for your target environment and allows you to move your schema, data, and uncontained objects from your source server to your target server.References:

<https://docs.microsoft.com/en-us/sql/dma/dma-overview>New QuestionA company plans to use Azure SQL Database to support a mission-critical application.The application must be highly available without performance degradation during maintenance windows.You need to implement the solution.Which three technologies should you implement? Each correct answer presents part of the solution.NOTE: Each correct selection is worth one point.A. Premium service tierB. Virtual machine Scale SetsC. Basic service tierD. SQL Data SyncE. Always On availability groupsF. Zone-redundant configurationAnswer: AEFExplanation:A: Premium/business critical service tier model that is based on a cluster of database engine processes. This architectural model relies on a fact that there is always a quorum of available database engine nodes and has minimal performance impact on your workload even during maintenance activities.E: In the premium model, Azure SQL database integrates compute and storage on the single node. High availability in this architectural model is achieved by replication of compute (SQL Server Database Engine process) and storage (locally attached SSD) deployed in 4-node cluster, using technology similar to SQL Server Always On Availability Groups. F: Zone redundant configurationBy default, the quorum-set replicas for the local storage configurations are created in the same datacenter. With the introduction of Azure Availability Zones, you have the ability to place the different replicas in the quorum-sets to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across

multiple zones as three gateway rings (GW).References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-high-availability>New QuestionA company plans to use Azure Storage for file storage purposes. Compliance rules require:- A single storage account to store all operations including reads, writes and deletes- Retention of an on-premises copy of historical operationsYou need to configure the storage account.Which two actions should you perform? Each correct answer presents part of the solution.NOTE: Each correct selection is worth one point.A.

Configure the storage account to log read, write and delete operations for service type BlobB. Use the AzCopy tool to download log data from \$logs/blobC. Configure the storage account to log read, write and delete operations for service-type tableD. Use the storage client to download log data from \$logs/tableE. Configure the storage account to log read, write and delete operations for service type queueAnswer: ABExplanation:Storage Logging logs request data in a set of blobs in a blob container named \$logs in your storage account. This container does not show up if you list all the blob containers in your account but you can see its contents if you access it directly.To view and analyze your log data, you should download the blobs that contain the log data you are interested in to a local machine. Many storage-browsing tools enable you to download blobs from your storage account; you can also use the Azure Storage team provided command-line Azure Copy Tool (AzCopy) to download your log data.References:

<https://docs.microsoft.com/en-us/rest/api/storageservices/enabling-storage-logging-and-accessing-log-data>1.|2019 Latest Braindump2go DP-200 Exam Dumps (PDF & VCE) Instant Download:<https://www.braindump2go.com/dp-200.html>2.|2019 Latest Braindump2go DP-200 Exam Questions & Answers Instant

Download:https://drive.google.com/drive/folders/1Lr1phOaAVcbL-_R5O-DFweNoShul0W13?usp=sharing