

[13/May/2019 Updated Pass DP-201 Exam By Training PassLeader Free DP-201 Braindumps

New Updated DP-201 Exam Questions from PassLeader DP-201 PDF dumps! Welcome to download the newest PassLeader DP-201 VCE dumps: <https://www.passleader.com/dp-201.html> (65 Q&As) Keywords: DP-201 exam dumps, DP-201 exam questions, DP-201 VCE dumps, DP-201 PDF dumps, DP-201 practice tests, DP-201 study guide, DP-201 braindumps, Designing an Azure Data Solution Exam P.S. New DP-201 dumps PDF:

<https://drive.google.com/open?id=1VdzP5HksyU93Arqn65qPe5UFEm2SxooH> P.S. New DP-100 dumps PDF:

<https://drive.google.com/open?id=1f70QWrCCtvNby8oY6BYvrMS16IXuRiR2> P.S. New DP-200 dumps PDF:

https://drive.google.com/open?id=1CTHwJ44u51T4tsb2qo8oThaQ5c_vwun1 NEW QUESTION 1 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. <https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview> NEW

QUESTION 2 You 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. Choose two.) A. Azure Cosmos DB B. Azure Service Bus C. Azure Storage Queue D. Apache Cassandra E. Kafka HDInsight Answer: A E Explanation: Option A: 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 applications.- Azure Cosmos DB change feed, which streams new data to the batch layer for HDInsight to process.- The Spark to Azure Cosmos DB Connector. Option E: You can use Apache Spark to stream data into or out of Apache Kafka on HDInsight using DStreams.

<https://docs.microsoft.com/en-us/azure/cosmos-db/lambda-architecture> NEW QUESTION 3 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 MLlib 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. MLlib 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

<https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-machine-learning-mllib-ipynb> NEW QUESTION 4 A company is developing a solution to manage inventory data for a group of automotive repair shops. The solution will use Azure SQL Data Warehouse as the data store. Shops will upload data every 10 days. Data corruption checks must run each time data is uploaded. If corruption is detected, the corrupted data must be removed. You need to ensure that upload processes and data corruption checks do not impact reporting and analytics processes that use the data warehouse. Solution: Insert data from shops and perform the data corruption check in a transaction. Rollback transfer if corruption is detected. Does the solution meet the goal? A. Yes B. No Answer: B Explanation: Instead, create a user-defined restore point before data is uploaded. Delete the restore point after data corruption checks complete. <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/backup-and-restore> NEW QUESTION

5 A company is developing a solution to manage inventory data for a group of automotive repair shops. The solution will use Azure SQL Data Warehouse as the data store. Shops will upload data every 10 days. Data corruption checks must run each time data is uploaded. If corruption is detected, the corrupted data must be removed. You need to ensure that upload processes and data corruption checks do not impact reporting and analytics processes that use the data warehouse. Solution: Create a user-defined restore point before data is uploaded. Delete the restore point after data corruption checks complete. Does the solution meet the goal? A. Yes B. No Answer: A Explanation: User-Defined Restore Points. This feature enables you to manually trigger snapshots to create restore points of your data warehouse before and after large modifications. This capability ensures that restore points are logically consistent, which provides additional data protection in case of any workload interruptions or user errors for quick recovery time. Note: A data warehouse restore is a new data warehouse that is created from a restore point of an existing or deleted data warehouse. Restoring your data warehouse is an essential part of any business continuity and disaster recovery strategy because it

re-creates your data after accidental corruption or deletion.

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/backup-and-restore> NEW QUESTION 6 You are designing an Azure SQL Data Warehouse. You plan to load millions of rows of data into the data warehouse each day. You must ensure that staging tables are optimized for data loading. You need to design the staging tables. What type of tables should you recommend? A. Round-robin distributed table B. Hash-distributed table C. Replicated table D. External table Answer: A Explanation: To achieve the fastest loading speed for moving data into a data warehouse table, load data into a staging table. Define the staging table as a heap and use round-robin for the distribution option. Incorrect: Not B: Consider that loading is usually a two-step process in which you first load to a staging table and then insert the data into a production data warehouse table. If the production table uses a hash distribution, the total time to load and insert might be faster if you define the staging table with the hash distribution. Loading to the staging table takes longer, but the second step of inserting the rows to the production table does not incur data movement across the distributions. <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data> NEW QUESTION 7 Drag and Drop You are designing a Spark job that performs batch processing of daily web log traffic. When you deploy the job in the production environment, it must meet the following requirements: - Run once a day. - Display status information on the company intranet as the job runs. You need to recommend technologies for triggering and monitoring jobs. Which technologies should you recommend? (To answer, drag the appropriate technologies to the correct locations. Each technology 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.) **Answer Area**

Technologies	Requirements
Livy	Triggering of jobs Monitoring of jobs
Beeline	
Azure Logic App	
Azure API App	

Answer: **Answer Area**

Technologies	Requirement	Technology
Livy	Triggering of jobs	Livy
Beeline	Monitoring of jobs	Beeline
Azure Logic App		
Azure API App		

www.passleader.com

Explanation: Box 1: Livy. You can use Livy to run interactive Spark shells or submit batch jobs to be run on Spark. Box 2: Beeline. Apache Beeline can be used to run Apache Hive queries on HDInsight. You can use Beeline with Apache Spark. Note: Beeline is a Hive client that is included on the head nodes of your HDInsight cluster. Beeline uses JDBC to connect to HiveServer2, a service hosted on your HDInsight cluster. You can also use Beeline to access Hive on HDInsight remotely over the internet.

<https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-livy-rest-interface>

<https://docs.microsoft.com/en-us/azure/hdinsight/hadoop/apache-hadoop-use-hive-beeline> NEW QUESTION 8 HotSpot 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.) **Answer Area**

Res

Answer: **Answer Area**

www.passleader.com

Restore order	Backup type
first	<div><div></div><div>full weekly backup</div><div>full daily backup</div><div>differential weekly backup</div><div>differential daily backup</div></div>
second	<div><div></div><div>full daily backup</div><div>differential backup from the last 12 hours</div><div>all differential backups since the last full backup</div><div>all log backups since the last full backup</div></div>
third	<div><div></div><div>all log backups since the last differential backup</div><div>differential backup from the last 12 hours</div><div>all differential backups since the last full backup</div><div>all log backups since the last full backup</div></div>

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.

<https://azure.microsoft.com/sv-se/blog/azure-sql-database-point-in-time-restore/> NEW QUESTION 9A company stores sensitive information about customers and employees in Azure SQL Database. You need to ensure that the sensitive data remains encrypted in transit and at rest. What should you recommend? A. Transparent Data EncryptionB. Always Encrypted with secure enclavesC. Azure Disk EncryptionD. SQL Server AlwaysOn Answer: BExplanation: Not A: Transparent Data Encryption (TDE) encrypts SQL Server, Azure SQL Database, and Azure SQL Data Warehouse data files, known as encrypting data at rest. TDE does not provide encryption across communication channels.

<https://cloudblogs.microsoft.com/sqlserver/2018/12/17/confidential-computing-using-always-encrypted-with-secure-enclaves-in-sql-server-2019-preview/> NEW QUESTION 10You plan to use Azure SQL Database to support a line of business app. You need to identify sensitive data that is stored in the database and monitor access to the data. Which three actions should you recommend?

(Each correct answer presents part of the solution. Choose three.) A. Enable Data Discovery and Classification.B. Implement Transparent Data Encryption (TDE).C. Enable Auditing.D. Run Vulnerability Assessment.E. Use Advanced Threat Protection.

Answer: CDE NEW QUESTION 11..... Case Study 1 - Trey ResearchTrey Research is a technology innovator. The company partners with regional transportation department office to build solutions that improve traffic flow and safety. The company is developing the following solutions:..... NEW QUESTION 41You need to design the vehicle images storage solution. What should you recommend? A. Azure Media ServicesB. Azure Premium Storage AccountC. Azure Redis CacheD. Azure Cosmos DB
Answer: BExplanation:Premium Storage stores data on the latest technology Solid State Drives (SSDs) whereas Standard Storage stores data on Hard Disk Drives (HDDs). Premium Storage is designed for Azure Virtual Machine workloads which require consistent high IO performance and low latency in order to host IO intensive workloads like OLTP, Big Data, and Data Warehousing on platforms like SQL Server, MongoDB, Cassandra, and others. With Premium Storage, more customers will be able to lift-and-shift demanding enterprise applications to the cloud.

<https://azure.microsoft.com/es-es/blog/introducing-premium-storage-high-performance-storage-for-azure-virtual-machine-workloads>

/ NEW QUESTION 42You need to recommend an Azure SQL Database pricing tier for Planning Assistance. Which pricing tier should you recommend? A. Business critical Azure SQL Database single database.B. General purpose Azure SQL Database Managed Instance.C. Business critical Azure SQL Database Managed Instance.D. General purpose Azure SQL Database single database. Answer: BExplanation:Azure resource costs must be minimized where possible. Data used for Planning Assistance must be stored in a sharded Azure SQL Database. The SLA for Planning Assistance is 70 percent, and multiday outages are permitted.

NEW QUESTION 43You need to design the runtime environment for the Real Time Response system. What should you recommend? A. General Purpose nodes without the Enterprise Security package.B. Memory Optimized nodes without the Enterprise Security package.C. Memory Optimized nodes with the Enterprise Security package.D. General Purpose nodes with the Enterprise Security package. Answer: B NEW QUESTION 44..... Case Study 2 - Proseware, Inc.Proseware, Inc. 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.The solution has the following technical requirements:..... NEW QUESTION 51You need to design the solution for analyzing customer data. What should you recommend? A. Azure DatabricksB. Azure Data Lake StorageC. Azure SQL Data WarehouseD. Azure Cognitive ServicesE.

Azure Batch Answer: AExplanation:Customer data must be analyzed using managed Spark clusters. You create spark clusters through Azure Databricks.<https://docs.microsoft.com/en-us/azure/azure-databricks/quickstart-create-databricks-workspace-portal>

NEW QUESTION 52You 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 Warehouse
Answer: CExplanation: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.<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blobs-introduction>

NEW QUESTION 53You need to design a backup solution for the processed customer data. What should you include in the design? A. AzCopyB. AdlCopyC. Geo-RedundancyD. Geo-Replication Answer: CExplanation: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.

<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy-grs> NEW QUESTION 54..... Case Study 3 - Contoso Contoso has the following virtual machines (VMs):..... NEW QUESTION 61You 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 DiskB. Premium SSD Managed DiskC. Ultra SSD Managed Disk Answer: CExplanation:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disks-types> NEW QUESTION 62You need to optimize storage for CONT_SQL3. What should you recommend? A. AlwaysOnB. Transactional processingC. GeneralD. Data warehousing
Answer: BExplanation: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

<https://www.databasejournal.com/features/mssql/overview-of-in-memory-technologies-of-azure-sql-database.html> NEW

QUESTION 63 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 Dataloader. Answer: B NEW QUESTION 64..... Download the newest PassLeader DP-201 dumps from passleader.com now! 100% Pass Guarantee! DP-201 PDF dumps & DP-201 VCE dumps: <https://www.passleader.com/dp-201.html> (65 Q&As) (New Questions Are 100% Available and Wrong Answers Have Been Corrected! Free VCE simulator!) P.S. New DP-201 dumps PDF:

<https://drive.google.com/open?id=1VdzP5HksyU93Arqn65qPe5UFEm2Sxooh> P.S. New DP-100 dumps PDF:

<https://drive.google.com/open?id=1f70QWrCCtvNby8oY6BYvrMS16IXuRiR2> P.S. New DP-200 dumps PDF:

https://drive.google.com/open?id=1CTHwJ44u5lT4tsb2qo8oThaQ5c_vwun1