

[23/Feb/2019 Updated PassLeader 103q 70-475 VCE Braindumps For Free Download

New Updated 70-475 Exam Questions from PassLeader 70-475 PDF dumps! Welcome to download the newest PassLeader 70-475 VCE dumps: <https://www.passleader.com/70-475.html> (103 Q&As) Keywords: 70-475 exam dumps, 70-475 exam questions, 70-475 VCE dumps, 70-475 PDF dumps, 70-475 practice tests, 70-475 study guide, 70-475 braindumps, Designing and Implementing Big Data Analytics Solutions Exam P.S. New 70-475 dumps PDF:

https://drive.google.com/open?id=1q2V0jQVJHIBUpXmZlXa_WGStspUYOQF0 P.S. New 70-473 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpMUdyRGNIS01WMjA NEW QUESTION 86 You plan to deploy a Microsoft Azure SQL data warehouse and a web application. The data warehouse will ingest 5 TB of data from an on-premises Microsoft SQL Server database daily. The web application will query the data warehouse. You need to design a solution to ingest data into the data warehouse. Solution: You use AzCopy to transfer the data as text files from SQL Server to Azure Blob storage, and then you use PolyBase to run Transact-SQL statements that refresh the data warehouse database. Does this meet the goal? A. Yes B. No Answer: A Explanation: If you need the best performance, then use PolyBase to import data into Azure SQL warehouse.

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-migrate-data> NEW QUESTION 87 You plan to deploy a Microsoft Azure SQL data warehouse and a web application. The data warehouse will ingest 5 TB of data from an on-premises Microsoft SQL Server database daily. The web application will query the data warehouse. You need to design a solution to ingest data into the data warehouse. Solution: You use the bcp utility to export CSV files from SQL Server and then to import the files to Azure SQL Data Warehouse. Does this meet the goal? A. Yes B. No Answer: B Explanation: If you need the best performance, then use PolyBase to import data into Azure SQL warehouse.

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-migrate-data> NEW QUESTION 88 You have an Apache Storm cluster. You need to ingest data from a Kafka queue. Which component should you use to consume data emitted from Kafka? A. a bolt B. a Microsoft Azure Service Bus queue C. a spout D. Flume Answer: C Explanation: To perform real-time computation on Storm, we create "topologies". A topology is a graph of a computation, containing a network of nodes called "Spouts" and "Bolts". In a Storm topology, a Spout is the source of data streams and a Bolt holds the business logic for analyzing and processing those streams. The org.apache.storm.kafka.KafkaSpout component reads data from Kafka.

<https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-apache-storm-with-kafka>

<https://hortonworks.com/blog/storm-kafka-together-real-time-data-refinery/> NEW QUESTION 89 You plan to use Microsoft Azure IoT Hub to capture data from medical devices that contain sensors. You need to ensure that each device has its own credentials. The solution must minimize the number of required privileges. Which policy should you apply to the devices? A. iothubowner B. service C. registryReadWrite D. device Answer: D Explanation: Per-Device Security Credentials. Each IoT Hub contains an identity registry. For each device in this identity registry, you can configure security credentials that grant DeviceConnect permissions scoped to the corresponding device endpoints. By the way, for option A: an iothubowner would have all permissions.

NEW QUESTION 90 You have a Microsoft Azure Data Factory that loads data to an analytics solution. You receive an alert that an error occurred during the last processing of a data stream. You debug the problem and solve an error. You need to process the data stream that caused the error. What should you do? A. From Azure Cloud Shell, run the az dla job command. B. From Azure Cloud Shell, run the az batch job enable command. C. From PowerShell, run the Resume-AzureRmDataFactoryPipeline cmdlet. D. From PowerShell, run the Set-AzureRmDataFactorySliceStatus cmdlet. Answer: D Explanation: ADF operates on data in batches known as slices. Slices are obtained by querying data over a date-time window, for example, a slice may contain data for a specific hour, day, or week. Such a query would use a date-time column notionally representing an effective date or a last-modified date.

<https://medium.com/@benhatton/how-to-construct-time-slices-of-your-data-for-azure-data-factory-4977c2a85cde> NEW

QUESTION 91 You have a Microsoft Azure Data Factory pipeline. You discover that the pipeline fails to execute because data is missing. You need to rerun the failure in the pipeline. Which cmdlet should you use? A. Set-AzureRmAutomationJob B. Set-AzureRmDataFactorySliceStatus C. Resume-AzureRmDataFactoryPipeline D. Resume-AzureRmAutomationJob Answer: B Explanation: Use some PowerShell to inspect the ADF activity for the missing file error. Then simply set the dataset slice to either skipped or ready using the cmdlet to override the status.

<https://stackoverflow.com/questions/42723269/azure-data-factory-pipelines-are-failing-when-no-files-available-in-the-source> NEW

QUESTION 92 You plan to analyze the execution logs of a pipeline to identify failures by using Microsoft Power BI. You need to automate the collection of monitoring data for the planned analysis. What should you do from Microsoft Azure? A. Create a Data Factory Set B. Save a Data Factory Log C. Add a Log Profile D. Create an Alert Rule Email Answer: A Explanation: You can

import the results of a Log Analytics log search into a Power BI dataset so you can take advantage of its features such as combining data from different sources and sharing reports on the web and mobile devices. To import data from a Log Analytics workspace into Power BI, you create a dataset in Power BI based on a log search query in Log Analytics. The query is run each time the dataset is refreshed. You can then build Power BI reports that use data from the dataset.

<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/powerbi> NEW QUESTION 93 You have an Apache Hive cluster in Microsoft Azure HDInsight. The cluster contains 10 million data files. You plan to archive the data. The data will be analyzed monthly. You need to recommend a solution to move and store the data. The solution must minimize how long it takes to move the data and must minimize costs. Which two services should you include in the recommendation? (Each correct answer presents part of the solution. Choose two.) A. Azure Queue storage B. Microsoft SQL Server Integration Services (SSIS) C. Azure Table Storage D. Azure Data Lake E. Azure Data Factory Answer: DE Explanation: D: To analyze data in HDInsight cluster, you can store the data either in Azure Storage, Azure Data Lake Storage Gen 1/Azure Data Lake Storage Gen 2, or both. Both storage options enable you to safely delete HDInsight clusters that are used for computation without losing user data. E: The Spark activity in a Data Factory pipeline executes a Spark program on your own or on-demand HDInsight cluster. It handles data transformation and the supported transformation activities. <https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hadoop-use-data-lake-store>

<https://docs.microsoft.com/en-us/azure/data-factory/transform-data-using-spark> NEW QUESTION 94 You are building an Azure Analysis Services cube. The source data for the cube is located on premises in a Microsoft SQL Server database. You need to ensure that the Azure Analysis Services service can access the source data. What should you deploy to your Azure subscription? A. Azure Data Factory B. A network gateway in Azure C. A data gateway in Azure D. A site-to-site VPN Answer: C Explanation: Connecting to on-premises data sources from and Azure AS server require an on-premises gateway.

<https://azure.microsoft.com/en-in/blog/on-premises-data-gateway-support-for-azure-analysis-services/> NEW QUESTION 95 You have a Microsoft Azure SQL database that contains Personally Identifiable Information (PII). To mitigate the PII risk, you need to ensure that data is encrypted while the data is at rest. The solution must minimize any changes to front-end applications. What should you use? A. Transport Layer Security (TLS) B. Transparent Data Encryption (TDE) C. A Shared Access Signature (SAS) D. The ENCRYPTBYPASSPHRASE T-SQL Function Answer: B Explanation: Transparent data encryption (TDE) helps protect Azure SQL Database, Azure SQL Managed Instance, and Azure Data Warehouse against the threat of malicious activity. It performs real-time encryption and decryption of the database, associated backups, and transaction log files at rest without requiring changes to the application. <https://docs.microsoft.com/en-us/azure/sql-database/transparent-data-encryption-azure-sql> NEW

QUESTION 96 A company named Fabrikam, Inc. has a web app hosted in Microsoft Azure. Millions of users visit the app daily. All of the user visits are logged in Azure Blob storage. Data analysts at Fabrikam built a dashboard that processes the user visit logs. Fabrikam plans to use an Apache Hadoop cluster on Azure HDInsight to process queries. The queries will access the data only once. You need to recommend a query execution strategy. What is the best to recommend using to achieve the goal? (More than one answer choice may achieve the goal. Select the BEST answer.) A. Load the text files to ORC files, and then run dashboard queries on the ORC files. B. Load the text files to sequence files, and then run dashboard queries on the sequence files. C. Run the queries on the text files directly. D. Load the text files to parquet files, and then run dashboard queries on the parquet files. Answer: B Explanation: File format versatility and Intelligent caching: Fast analytics on Hadoop have always come with one big catch: they require up-front conversion to a columnar format like ORCFile, Parquet or Avro, which is time-consuming, complex and limits your agility. With Interactive Query Dynamic Text Cache, which converts CSV or JSON data into optimized in-memory format on-the-fly, caching is dynamic, so the queries determine what data is cached. After text data is cached, analytics run just as fast as if you had converted it to specific file formats.

<https://azure.microsoft.com/en-us/blog/azure-hdinsight-interactive-query-simplifying-big-data-analytics-architecture-and-operations/> NEW QUESTION 97 Hotspot You plan to implement a Microsoft Azure Data Factory pipeline. The pipeline will have custom business logic that requires a custom processing step. You need to implement the custom processing step by using C#. Which interface and method should you use? (To answer, select the appropriate options in the answer area.)

Answer Area

Interface:

	▼
ICustomActivity	
IDotNetActivity	
IGenericActivity	

www.passleader.com

Method:

	▼
Copy	
Execute	
Run	
Update	

Answer: **Answer Area**

Interface:

	▼
ICustomActivity	
IDotNetActivity	
IGenericActivity	

www.passleader.com

Method:

	▼
Copy	
Execute	
Run	
Update	

Explanation:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/data-factory/v1/data-factory-use-custom-activities.md> NEW QUESTION 98 Drag and Drop You have an analytics solution in Microsoft Azure that must be operationalized. You have the relevant data in Azure Blob storage. You use an Azure HDInsight Cluster to process the data. You plan to process the raw data files by using Azure HDInsight. Azure Data Factory will operationalize the solution. You need to create a data factory to orchestrate the data movement. Output data must be written back to Azure Blob storage. Which four actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Actions

Answer Area

Create input and output datasets for the files in Azure.

1

Rename the data factory hub.

2

Create a data factory.

3

Create a client to connect to Azure Blob storage.

4

Create linked services for Azure Blob storage and the Azure HDInsight cluster.

Create an Azure HDInsight activity in a pipeline to process the data.

Answer:

Actions

Answer Area

Rename the data factory hub.

1

Create a data factory.

2

Create linked services for Azure Blob storage and the Azure HDInsight cluster.

3

Create input and output datasets for the files in Azure.

4

Create an Azure HDInsight activity in a pipeline to process the data.

Create a client to connect to Azure Blob storage.

www.passleader.com

NEW QUESTION 99 Drag and Drop You need to implement a security solution for Microsoft Azure SQL database. The solution must meet the following requirements:- Ensure that users can see the data from their respective department only.- Prevent administrators from viewing the data. Which feature should you use for each requirement? (To answer, drag the appropriate features to the correct requirements. Each feature 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.)

Features

Answer Area

Always Encrypted

Dynamic Data Masking

Row-level security (RLS)

Transparent Data Encryption (TDE)

Ensure that users can see the data from their respective department only:

www.passleader.com

Prevent administrators from viewing the data:

Answer: **Features**

- Row-level security (RLS)
- Transparent Data Encryption (TDE)

Answer Area

Ensure that users can see the data from their respective department only:

www.passleader.com

Prevent administrators from viewing the data:

Dynamic Data Masking

Always Encrypted

NEW QUESTION 100 You deploy a Microsoft Azure SQL database. You create a job to upload customer data to the database. You discover that the job cannot connect to the database and fails. You verify that the database runs successfully in Azure. You need to run the job successfully. What should you create? A. a virtual network rule B. a network security group (NSG) C. a firewall rule D. a virtual network Answer: C Explanation: If the application persistently fails to connect to Azure SQL Database, it usually indicates an issue with one of the following: - Firewall configuration. The Azure SQL database or client-side firewall is blocking connections to Azure SQL Database. - Network reconfiguration on the client side: for example, a new IP address or a proxy server. - User error: for example, mistyped connection parameters, such as the server name in the connection string.

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-troubleshoot-common-connection-issues> NEW QUESTION 101

..... Download the newest PassLeader 70-475 dumps from passleader.com now! 100% Pass Guarantee! 70-475 PDF dumps & 70-475 VCE dumps: <https://www.passleader.com/70-475.html> (103 Q&As) (New Questions Are 100% Available and Wrong Answers Have Been Corrected! Free VCE simulator!) P.S. New 70-475 dumps PDF:

https://drive.google.com/open?id=1q2V0jQVJHIBUpXmZIXa_WGStpUYOQF0 P.S. New 70-473 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpMUdyRGNIS01WMjA