

[2018-New-Exam100% Real Exam Questions-Braindump2go 70-775 Exam VCE and 70-775 PDF Dumps 40Q Download[11-21

2018 Feb New Microsoft 70-775 Exam Dumps with PDF and VCE Free Released Today! Following are some new 70-775 EReal Exam Questions:1.[2018 New 70-775 Exam Dumps (PDF & VCE) 40Q&As Download:

<https://www.braindump2go.com/70-775.html>2.[2018 New 70-775 Exam Questions & Answers Download:

https://drive.google.com/drive/folders/1uJ5XlxxFUMLKFts7UH_t9bHmKDKxIWDT?usp=sharingQuestion No : 11You plan to copy data from Azure Blob storage to an Azure SQL database by using Azure Data Factory.Which file formats can you use?A.

binary, JSON, Apache Parquet, and ORCB. OXPS, binary, text and JSONC. XML, Apache Avro, text, and ORCD. text, JSON, Apache Avro, and Apache ParquetAnswer: D Explanation:References:

<https://docs.microsoft.com/en-us/azure/data-factory/supported-file-formats-and-compression-codecs>Question No : 13Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.You are planning a big data infrastructure by using an Apache Spark cluster in Azure HDInsight. The cluster has 24 processor cores and 512 GB of memory.The architecture of the infrastructure is shown in the exhibit. (Click the Exhibit button.)The architecture will be used by the following users:The data sources in the batch layer share a common storage container. The following data sources are used:The business analysts report that they experience performance issues when they run the monitoring queries.You troubleshoot the performance issues and discover that the intermediate tables generated when the analysts run the queries cause pressure for the Java Virtual Machine (JVM) garbage collection per job.Which configuration settings should you modify to alleviate the performance issues?A. spark.sql.inMemoryColumnarStorage.batchSizeB. spark.sql.broadcastTimeoutC. spark.sql.files.openCostInBytesD. spark.sql.shuffle.partitionsAnswer: DQuestion No : 14Note: 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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.You are building a security tracking solution in Apache Kafka to parse security logs. The security logs record an entry each time a user attempts to access an application. Each log entry contains the IP address used to make the attempt and the country from which the attempt originated.You need to receive notifications when an IP address from outside of the United States is used to access the application.Solution: Create new topics. Create a file import process to send messages. Start the consumer and run the producer.Does this meet the goal?A. YesB. NoAnswer: AQuestion No : 15You use YARN to manage the resources for a Spark Thrift Server running on a Linux- based Apache Spark cluster in Azure HDInsight.You discover that the cluster does not fully utilize the resources. You want to increase resource allocation.You need to increase the number of executors and the allocation of memory to the Spark Thrift Server driver.Which two parameters should you modify? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.A. spark.dynamicAllocation.maxExecutorsB. spark.cores.maxC. spark.executor.memoryD. spark_thrift_cmd_optsE. spark.executor.instancesAnswer: A,C Explanation:References: <https://stackoverflow.com/questions/37871194/how-to-tune-spark-executor-number-cores-and-executor-memory>Question No : 16You have an Apache Hadoop cluster in Azure HDInsight that has a head node and three data nodes. You have a MapReduce job.You receive a notification that a data node failed.You need to identify which component cause the failure. Which tool should you use?A. JobTrackerB. TaskTrackerC. ResourceManagerD. ApplicationMasterAnswer: C Question No : 17HOTSPOTYou install the Microsoft Hive ODBC Driver on a computer that runs Windows 10 and has the 64-bit version of Microsoft Office 2016 installed.You deploy a new Apache Interactive Hive cluster in Azure HDInsight. The cluster is hosted at myHDICluster.azurehdinsight.net and contains a Hive table name hivesampletable that has 200,000 rows.You plan to use HiveQL exclusively for the queries. The queries will return from 6,000 to 10,000 rows 90 percent of the time.You need to configure a data source to ensure that you can use Microsoft Excel to access the data. The solution must ensure that the Hive queries execute as quickly as possible.How should you configure the Advanced Options from the Microsoft Hive ODBC Driver DSN Setup dialog box? To answer, select the appropriate options in the answer area.NOTE: Each correct selection is worth one point.Answer: Explanation:References: <https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-connect-excel-hive-odbc-driver>Question No : 18DRAG DROPYou have an Apache HBase cluster in Azure HDInsight. The cluster has a table named sales that contains a column family named customerfamily.You need to add a new column family named customeraddr to the sales table.How should you complete the command? To answer, drag the appropriate values to the correct targets. Each value 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. Answer: Explanation: Hbase > disable 'sales' Hbase > alter 'sales' ?customerfamily',{NAME => 'customeraddr', IN_MEMORY => false}, Hbase > enable 'sales' Question No : 19 Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question. You need to deploy an enterprise data warehouse that will support in-memory analytics. The data warehouse must support connections that use the Microsoft Hive ODBC Driver and Beeline. The data warehouse will be managed by using Apache Amrabi only. What should you do? A. Use an Azure PowerShell script to create and configure a premium HDInsight cluster. Specify Apache Hadoop as the cluster type and use Linux as the operating system. B. Use the Azure portal to create a standard HDInsight cluster. Specify Apache Spark as the cluster type and use Linux as the operating system. C. Use an Azure PowerShell script to create a standard HDInsight cluster. Specify Apache HBase as the cluster type and use Windows as the operating system. D. Use an Azure PowerShell script to create a standard HDInsight cluster. Specify Apache Storm as the cluster type and use Windows as the operating system. E. Use an Azure PowerShell script to create a premium HDInsight cluster. Specify Apache HBase as the cluster type and use Linux as the operating system. F. Use an Azure portal to create a standard HDInsight cluster. Specify Apache Interactive Hive as the cluster type and use Linux as the operating system. G. Use an Azure portal to create a standard HDInsight cluster. Specify Apache HBase as the cluster type and use Linux as the operating system. Answer: F Explanation: References:

<https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hadoop-use-interactive-hive> Question No : 20 Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question. You are implementing a batch processing solution by using Azure HDInsight. You have data stored in Azure. You need to ensure that you can access the data by using Azure Active Directory (Azure AD) identities. What should you do? A. Use a shuffle join in an Apache Hive query that stores the data in a JSON format. B. Use a broadcast join in an Apache Hive query that stores the data in an ORC format. C. Increase the number of spark.executor.cores in an Apache Spark job that stores the data in a text format. D. Increase the number of spark.executor.instances in an Apache Spark job that stores the data in a text format. E. Decrease the level of parallelism in an Apache Spark job that stores the data in a text format. F. Use an action in an Apache Oozie workflow that stores the data in a text format. G. Use an Azure Data Factory linked service that stores the data in Azure Data Lake. H. Use an Azure Data Factory linked service that stores the data in an Azure DocumentDB database. Answer: G Explanation: References:

<https://docs.microsoft.com/en-us/azure/data-factory/concepts-datasets-linked-services> Question No : 21 Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question. You need to deploy a NoSQL database to an HDInsight cluster. You will manage the server that host the database by using Remote Desktop. The database must use the key/value pair format in a columnar model. What should you do? A. Use an Azure PowerShell script to create and configure a premium HDInsight cluster. Specify Apache Hadoop as the cluster type and use Linux as the operating system. B. Use the Azure portal to create a standard HDInsight cluster. Specify Apache Spark as the cluster type and use Linux as the operating system. C. Use an Azure PowerShell script to create a standard HDInsight cluster. Specify Apache HBase as the cluster type and use Windows as the operating system. D. Use an Azure PowerShell script to create a standard HDInsight cluster. Specify Apache Storm as the cluster type and use Windows as the operating system. E. Use an Azure PowerShell script to create a premium HDInsight cluster. Specify Apache HBase as the cluster type and use Linux as the operating system. F. Use an Azure portal to create a standard HDInsight cluster. Specify Apache Interactive Hive as the cluster type and use Linux as the operating system. G. Use an Azure portal to create a standard HDInsight cluster. Specify Apache HBase as the cluster type and use Linux as the operating system. Answer: G Explanation: References:

<https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hbase-overview>!!!RECOMMEND!!!1. |2018 New 70-775 Exam Dumps (PDF & VCE) 40Q&As Download: <https://www.braindump2go.com/70-775.html> 2. |2018 New 70-775 Study Guide Video: YouTube Video: [YouTube.com/watch?v=FcygmKE0bd4](https://www.youtube.com/watch?v=FcygmKE0bd4)