

Microsoft

Exam 70-775

Perform Data Engineering on Microsoft Azure HDInsight

Version: 8.0

[Free Questions]

Question No : 33

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:

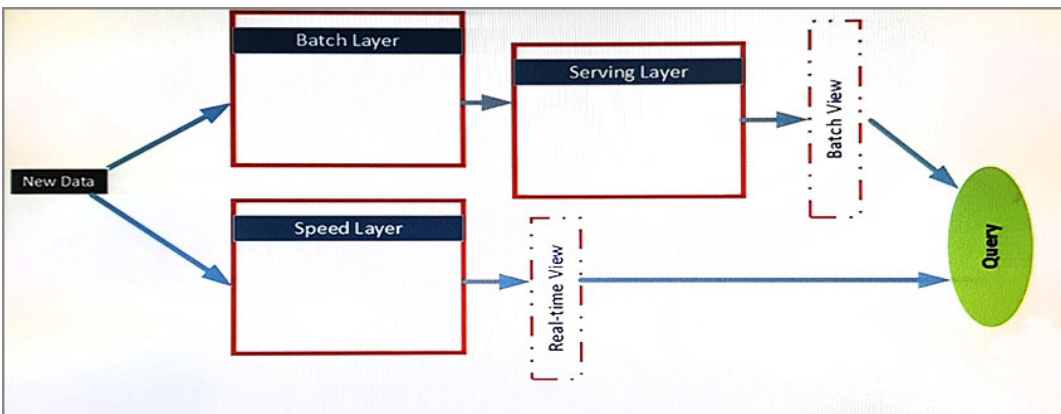
References: <https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hbase-overview>

Question No : 34

Note: 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.batchSize
- B. spark.sql.broadcastTimeout
- C. spark.sql.files.openCostInBytes
- D. spark.sql.shuffle.partitions

Answer: D

Question No : 35

You have an Azure HDInsight cluster.

You need to store data in a file format that maximizes compression and increases read performance.

Which type of file format should you use?

- A. ORC
- B. Apache Parquet
- C. Apache Avro
- D. Apache Sequence

Answer: A

Explanation:

References: <http://www.semantikoz.com/blog/orc-intelligent-big-data-file-format-hadoop-hive/>