

100% Money Back
Guarantee

Vendor:IBM

Exam Code:C2090-101

Exam Name:IBM Big Data Engineer

Version:Demo

QUESTION 1

Which of the following is the most effective method for improving query performance on large Hive tables?

- A. Indexing
- B. Bucketing
- C. Partitioning
- D. De-normalizing data

Correct Answer: B

Reference: <https://dzone.com/articles/how-to-improve-hive-query-performance-with-hadoop>

QUESTION 2

One option to import data to your BigInsights cluster is to use Sqoop.

Which of the following examples shows a data exchange with a DB2 database by using the built-in DB2 connector?

- A. `$SQOOP_HOME/bin/sqoop export --connect jdbc:db2://db2.my.com:50000/myDB --username db2user --password db2pwd --table db2tbl --export-dir /sqoop/dataFile.csv`
- B. `$SQOOP_HOME/bin/sqoop import -connect jdbc:db2://db2.my.com:50000/testdb -username db2user password db2pwd -table db2tbl -split-by tbl_primarykey -target-dir sqoopimports`
- C. `$SQOOP_HOME/bin/sqoop import --connect jdbc:db2://db2.my.com:50000/testdb --username db2user --password db2pwd --table db2tbl --split-by tbl_primarykey --target-dir sqoopimports`
- D. `$BIGINSIGHTS_HOME/bin/sqoop import --connect jdbc:db2://db2.my.com:50000/testdb -- username db2user --password db2pwd --table db2tbl --split-by tbl_primarykey --target-dir sqoopimports`

Correct Answer: C

Reference: https://www.ibm.com/support/knowledgecenter/SSPT3X_3.0.0/com.ibm.swg.im.infosphere.biginsights.import.doc/doc/data_warehouse_sqoop.html

QUESTION 3

Which of the following statement is TRUE with BigSheets?

- A. You can create any type of sheet from a parent workbook
- B. You must create a child workbook in order to create a chart
- C. You can delete a parent workbook without deleting the child workbooks
- D. You must run the workbook on the data to get the full results of the analysis

Correct Answer: A

Reference: https://www.ibm.com/support/knowledgecenter/SSPT3X_4.1.0/com.ibm.swg.im.infosphere.biginsights.analyze.doc/doc/bigsheets_con_workbooks.html

QUESTION 4

Bloom Filter in HBase can be used to determine which of the following?

- A. Whether a record exists in a region server
- B. Whether a record does not exist in a region server
- C. Items in the catalog database
- D. None of the above

Correct Answer: D

Reference: https://www.ibm.com/support/knowledgecenter/en/SSPT3X_3.0.0/com.ibm.swg.im.infosphere.biginsights.bigsql.doc/doc/bsql_create_hbase_table.html

QUESTION 5

Which of the following statements is TRUE regarding search visualization with Apache Hue?

- A. Hue utilizes Java libraries which must be installed on your system
- B. The Hue Beeswax application enables you to perform queries on HBase
- C. For optimal performance, the Hue Server should be one of the nodes within your Hadoop cluster
- D. The Hue Sqoop UI allows transferring data from a relational database to Hadoop, but not from Hadoop to a relational database

Correct Answer: D

Reference: <https://gethue.com/move-data-in-out-your-hadoop-cluster-with-the-sqoop/>

QUESTION 6

Which of the following statements is TRUE regarding search visualization with Apache Hue?

- A. Hue submits MapReduce jobs to Oozie
- B. No additional setup is required to secure your session cookies
- C. Hue applications require some code to be installed on the client
- D. The File Browser application allows you to perform keyword searches across your Hadoop data

Correct Answer: A

QUESTION 7

Setting HDFS folder permissions most directly relates to which of these PCI compliance requirements?

- A. Protect stored data
- B. Install and maintain a firewall
- C. Encrypt the transmission of sensitive data
- D. Assign a unique ID to each person with access to data

Correct Answer: C

QUESTION 8

Which ONE of the following statements regarding Sqoop is TRUE?

- A. HBase is not supported as an import target
- B. Data imported using Sqoop is always written to a single Hive partition
- C. Sqoop can be used to retrieve rows newer than some previously imported set of rows
- D. Sqoop can only append new rows to a database table when exporting back to a database

Correct Answer: C

Reference: <https://sqoop.apache.org/docs/1.4.1-incubating/SqoopUserGuide.html>

QUESTION 9

What are the key elements that IBM Big Match Probabilistic Matching Engine leverages?

- A. Wildcard and sorting
- B. Wildcard and phonetics
- C. Phonetics and searching
- D. Phonetics and nicknames

Correct Answer: D

Reference: https://www.ibm.com/support/knowledgecenter/SSWSR9_11.6.0/com.ibm.mdmhs.txn.ref.doc/r_searchPersonProbabilistic.html

QUESTION 10

Which of the following must happen before the Big SQL EXPLAIN command can execute?

- A. Run the ANALYZE command
- B. Set the COMPATIBILITY_MODE global variable
- C. Execute the SET HADOOP PROPERTY command
- D. Call the SYSPROC.SYSINSTALLOBJECTS procedure

Correct Answer: D

Reference: <https://developer.ibm.com/hadoop/docs/biginsights-ibm-open-platform/getting-started/tutorials/big-sql-hadoop-tutorial/getting-started-with-big-sql-4-0-lab-6-understanding-and-influencing-data-plans/>

QUESTION 11

Considering the following properties:

1.
Automated creation of target database schema and bulk extract and load
2.
Real-time replication subscriptions (with CDC)
3.
Managed workload for optimized performance of potentially thousands of artifacts
4.
Ensured governance around both data access as well as for metadata capture (to support data lineage and impact assessment)

Which tool supports all of the above?

- A. Pig
- B. JAQL
- C. Data Click
- D. BigSheets

Correct Answer: A

Reference: http://meta7.forsythe.com/_wss/clients/508/news_feed/20151130202508951.pdf

QUESTION 12

Which parameters are considered when configuring Big Match algorithm?

- A. Search and custom requirements
- B. Accuracy, search, and performance
- C. Adaptive weighting and standardization
- D. Empirical components, accuracy, and performance

Correct Answer: B