

DUMPSBOSS.

IBM Big Data Engineer

IBM C2090-101

Version Demo

Total Demo Questions: 10

Total Premium Questions: 106

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co

dumpsboss.co

QUESTION NO: 1

Which of the following Hadoop command flags are used to copy data to the HDFS? (Choose two.)

- A. cp
- B. put
- C. get
- D. copyToLocal
- E. copyFromLocal

ANSWER: D E

Explanation:

Reference: https://www.ibm.com/support/knowledgecenter/SSPT3X_4.2.0/com.ibm.swg.im.infosphere.biginsights.import.doc/doc/data_rest_shell.html

QUESTION NO: 2

Which of the following statements regarding Sqoop is TRUE? (Choose two.)

- A. All columns in a table must be imported
- B. Sqoop bypasses MapReduce for enhanced performance
- C. Each row from a source table is represented as a separate record in HDFS
- D. When using a password file, the file containing the password must reside in HDFS
- E. Multiple options files can be specified when invoking Sqoop from the command line

ANSWER: C E

Explanation:

Reference: <https://data-flair.training/blogs/apache-sqoop-tutorial/>

QUESTION NO: 3

Which file format support Column data compression? (Choose two.)

- A. Text
- B. Avro
- C. RCFile
- D. Parquet
- E. Sequence_text

ANSWER: C D

QUESTION NO: 4

Considering the following properties:

- Automated creation of target database schema and bulk extract and load
- Real-time replication subscriptions (with CDC)
- Managed workload for optimized performance of potentially thousands of artifacts
- 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

ANSWER: A

Explanation:

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

QUESTION NO: 5

Which of the following is NOT a characteristic of a consistent region?

- A. Support for non-deterministic applications
- B. Support for single-threaded operators only
- C. Allows “at-least-once” processing by using replayable data sources

D. Allows “exactly-once” processing if using an external system that can detect duplicate tuples or can restore its state after a failure

ANSWER: B

Explanation:

Reference: <https://developer.ibm.com/streamsdev/2015/02/20/processing-tuples-least-infospherestreams-consistent-regions/>

QUESTION NO: 6

Which of the following is not a data-processing operations that is supported in Pig Latin?

- A. filter
- B. joins
- C. group by
- D. logistic regression

ANSWER: D

Explanation:

Reference: <https://pig.apache.org/docs/r0.15.0/basic.html>

QUESTION NO: 7

When we create a new table in Hive, which clause can be used in HiveSQL to indicate the storage file format?

- A. SAVE AS
- B. MAKE AS
- C. FORMAT AS
- D. STORED AS

ANSWER: D

QUESTION NO: 8

Which of the following statements are TRUE regarding the use of Data Click to load data into BigInsights? (Choose two.)

- A. Big SQL cannot be used to access the data moved in by Data Click because the data is in Hive

- B. You must import metadata for all sources and targets that you want to make available for Data Click activities
- C. Connections from the relational database source to HDFS are discovered automatically from within Data Click
- D. Hive tables are automatically created every time you run an activity that moves data from a relational database into HDFS
- E. HBase tables are automatically created every time you ran an activity that moves data from a relational database into HDFS

ANSWER: C E

Explanation:

Reference: https://www.ibm.com/support/knowledgecenter/en/SSZJPZ_11.3.0/com.ibm.swg.im.iis.dataclick.doc/topics/hivetables.html

QUESTION NO: 9

Which of the following are CRUD operations available in HBase? (Choose two.)

- A. HTable.Put
- B. HTable.Read
- C. HTable.Delete
- D. HTable.Update
- E. HTable.Remove

ANSWER: A C

Explanation:

Reference: https://www.tutorialspoint.com/hbase/hbase_client_api.htm

QUESTION NO: 10

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 builtin 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

ANSWER: C

Explanation:

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