

C2090-101^{Q&As}

IBM Big Data Engineer

Pass IBM C2090-101 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass2lead.com/c2090-101.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by IBM Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



QUESTION 1

Which statement is TRUE when loading data into Hadoop and creating Big SQL tables?

- A. It is optional to have INSERT privileges granted to LOAD into a table with the APPEND option
- B. You can either have INSERT or DELETE privileges granted to LOAD into a table with the OVERWRITE
- C. Authentication information is not necessary when you connect to a secured InfoSphere BigInsights cluster
- D. By using the LOAD HADOOP USING command, you can import data from external data sources into target Big SQL tables

Correct Answer: D

Reference: https://www.ibm.com/support/knowledgecenter/en/SSCRJT_5.0.1/com.ibm.swg.im.bigsql.db2biga.doc/doc/biga_load_from.html

QUESTION 2

Which keyword can be inserted before a Jaql statement to understand how it may be rewritten into MapReduce?

- A. explain
- B. describe
- C. transcribe
- D. mapreduce

Correct Answer: D

QUESTION 3

Which BigInsights components are essential for Big Match operations?

- A. Social Data Analytics (SDA)
- B. HDFS, HBase, and BigSheets
- C. HDFS and HBase
- D. MapReduce framework, BigInsights cluster management

Correct Answer: D

Reference: https://www.ibm.com/support/knowledgecenter/SSWSR9_11.3.0/com.ibm.swg.im.mdmhs.pmebi.doc/topics/pme_bi_architecture.html

QUESTION 4

Which of the following is most commonly used by Hadoop to move data between clusters?

- A. Pig
- B. FTP
- C. JAQL
- D. distcp

Correct Answer: D

Reference: <https://developer.ibm.com/hadoop/2016/02/05/fast-can-data-transferred-hadoop-clustersusing-distcp/>

QUESTION 5

Which annotation is used to exclude operators from a consistent region?

- A. @consistent
- B. @standalone
- C. @independent
- D. @autonomous

Correct Answer: D

Reference: https://www.ibm.com/support/knowledgecenter/en/SSCRJU_4.2.1/com.ibm.streams.dev.doc/doc/consistentscope.html

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

Which ONE of the following statements regarding Sqoop is TRUE?

- A. By default, data is compressed with Sqoop
- B. Sqoop can only read committed transactions from a source database, not uncommitted ones
- C. When performing parallel imports, Sqoop always uses the primary key column in a table as the splitting column
- D. When performing parallel imports, each degree of parallelism corresponds to a concurrent database connection

Correct Answer: B

By default, Sqoop uses the read committed transaction isolation in the mappers to import data. This may not be the ideal in all ETL workflows and it may be desired to reduce the isolation guarantees. The `--relaxedisolation` option can be used to instruct Sqoop to use read uncommitted isolation level.

Reference: <https://sqoop.apache.org/docs/1.4.6/SqoopUserGuide.html>

QUESTION 8

When you configure a MapReduce job, the inputs can include:

- A. A single file
- B. Paths to one or more directories
- C. A file pattern (e.g., `my_path/*.csv`)
- D. All of the above

Correct Answer: C

QUESTION 9

In which of the following the Flume open-source software is written?

- A. C / C++
- B. Primarily Java
- C. Python and related tools
- D. A combination of C++ and Perl

Correct Answer: B

Reference: https://en.wikipedia.org/wiki/Apache_Flume

QUESTION 10

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 11

Which of the following techniques is NOT employed by Big SQL to improve performance?

- A. Query Optimization
- B. Predicate Push down
- C. Compression efficiency
- D. Load data into DB2 and return the data

Correct Answer: A

Reference: https://www.ibm.com/support/knowledgecenter/en/SSZLC2_7.0.0/com.ibm.commerce.developer.soa.doc/refs/rsdperformanceworkspaces.htm

QUESTION 12

Which of the following statements is TRUE regarding the use of Data Click to load data into BigInsights?

- A. Data Click uses Map Reduce to create a Hive table
- B. Only relational databases are supported as sources for Data Click
- C. Only whole tables can be moved from a relational database source, not individual columns within the table
- D. If the metadata that you specify in InfoSphere Data Click activities changes, you need to reimport the metadata by using InfoSphere Metadata Asset Manager

Correct Answer: D

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

QUESTION 13

Given a file named readme.txt, which command will copy the readme.txt file to the directory on the HDFS?

- A. `hadoop fs -cp readme.txt hdfs://test.ibm.com:9000/`

- B. `hadoop fs -cp hdfs://test.ibm.com:9000/ readme.txt`
- C. `hadoop fs -put readme.txt hdfs://test.ibm.com:9000/`
- D. `hadoop fs -put hdfs://test.ibm.com:9000/ readme.txt`

Correct Answer: B

QUESTION 14

Which of the following is TRUE regarding a schema In Jaql?

- A. It specifies a particular type for each field
- B. It is used only when Jaql includes SQL statements
- C. It must be provided when the data format is not otherwise known
- D. It can be used both as a data constraint and to optimize processing and storage

Correct Answer: B

QUESTION 15

Which statement is TRUE about Big SQL?

- A. The table definition can include other attributes such as the primary key or check constraints
- B. When using Big SQL, the CREATE TABLE statement cannot be embedded in an application program
- C. If a sub-table is being defined, the authorization ID can be either the same as the owner of the root table or an equivalent
- D. When defining a staging table associated with a materialized query table, the privileges held by the authorization ID of the statement only works with DBADM authority

Correct Answer: D

Reference: https://www.ibm.com/support/knowledgecenter/fi/SSCRJT_5.0.2/com.ibm.swg.im.bigsql.commsql.doc/doc/r0000927.html

[C2090-101 PDF Dumps](#)

[C2090-101 Practice Test](#)

[C2090-101 Braindumps](#)