



# C2090-101<sup>Q&As</sup>

IBM Big Data Engineer

**Pass IBM C2090-101 Exam with 100% Guarantee**

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

<https://www.pass4itsure.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 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 2

In order for an SPSS Modeler stream to be incorporated for use in an InfoSphere Streams application leveraging SPSS Modeler Solution Publisher, you need to:

- A. add a Type node
- B. insert any Output node
- C. add a Table node as the terminal node
- D. Make the terminal node a scoring branch

Correct Answer: D

---

### QUESTION 3

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

Correct Answer: AC

Reference: [https://www.tutorialspoint.com/hbase/hbase\\_client\\_api.htm](https://www.tutorialspoint.com/hbase/hbase_client_api.htm)

---

**QUESTION 4**

When embedding SPSS models within InfoSphere Streams, what SPSS product must be installed on the same machine with InfoSphere Streams?

- A. SPSS Modeler
- B. SPSS Solution Publisher
- C. SPSS Accelerator for InfoSphere Streams
- D. None, the SPSS software runs remotely to the Streams machine

Correct Answer: B

---

**QUESTION 5**

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 6**

Which statement is TRUE concerning optimizing the load performance?

- A. You can improve the performance by increasing the number of map tasks assigned to the load
- B. When loading large files the number of files that you load does not impact the performance of the LOAD HADOOP statement
- C. You can improve the performance by decreasing the number of map tasks that are assigned to the load and adjusting the heap size
- D. It is advantageous to run the LOAD HADOOP statement directly pointing to large files located in the host file system as opposed to copying the files to the DFS prior to load

Correct Answer: A

Reference: [https://www.ibm.com/support/knowledgecenter/en/SSCRJT\\_5.0.3/com.ibm.swg.im.bigsql.doc/doc/bigsql\\_loadperf.html](https://www.ibm.com/support/knowledgecenter/en/SSCRJT_5.0.3/com.ibm.swg.im.bigsql.doc/doc/bigsql_loadperf.html)

---

**QUESTION 7**

Which of the following statements regarding importing streaming data from InfoSphere Streams into Hadoop is TRUE?

- A. InfoSphere Streams utilizes Flume to interface to Hadoop
- B. The HDFSFileSink operator writes files in parallel to a Hadoop Distributed File System
- C. Buffering techniques are used to process incoming streams from InfoSphere Streams
- D. When you use the HDFS operators to access GPFS, you must install InfoSphere Streams on an InfoSphere Big Insights data node

Correct Answer: C

Reference: <https://books.google.com.pk/books?id=JWfRAGAAQBAJandpg=PA147andlpg=PA147anddq=IBM+Buffering+techniques+are+used+to+process+incoming+streams+from+InfoSphere+Streamsandsource=blandots=Z6XhA0Owkandsig=ACfU3U3T9ydrZHWTMvB31qQOyf6FtoDQgwandhl=enandsa=Xandved=2ahUKEwj6wsnvqvfAhVUSxUIHWNDAYcQ6AEwA3oECBEQAQ#v=onepageandq=IBM%20Buffering%20techniques%20are%20used%20to%20process%20incoming%20streams%20from%20InfoSphere%20Streamsandf=false>

---

**QUESTION 8**

The Distributed File Copy application copies files to and from a remote source to the InfoSphere BigInsights distributed file system by using data movement options.

Which of the following is the complete list of options for moving data with Distributed File Copy?

- A. File Transfer Protocol (FTP) and SFTP
- B. Hadoop Distributed File System (HDFS) and FTP
- C. General Parallel File System (GPFS), FTP, and SFTP
- D. Hadoop Distributed File System (HDFS), FTP, SFTP, GPFS, and also copy files to and from your local file system

Correct Answer: D

Reference: [https://www.ibm.com/support/knowledgecenter/en/SSPT3X\\_3.0.0/com.ibm.svg.im.infosphere.biginsights.tut.doc/doc/tut\\_Less\\_imp\\_distcopy.html](https://www.ibm.com/support/knowledgecenter/en/SSPT3X_3.0.0/com.ibm.svg.im.infosphere.biginsights.tut.doc/doc/tut_Less_imp_distcopy.html)

---

**QUESTION 9**

Which Big SQL statement can be used to store a single row into an HBase table?

- A. LOAD HIVE DATA
- B. INSERT INTO HBASE
- C. CREATE HBASE TABLE
- D. LOAD USING ... INTO HBASE TABLE



Correct Answer: B

Reference: [https://www.ibm.com/support/knowledgecenter/en/SSPT3X\\_3.0.0/com.ibm.swg.im.infosphere.biginsights.bigsql.doc/doc/bsql\\_insert.html](https://www.ibm.com/support/knowledgecenter/en/SSPT3X_3.0.0/com.ibm.swg.im.infosphere.biginsights.bigsql.doc/doc/bsql_insert.html)

---

#### QUESTION 10

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

Correct Answer: D

---

#### QUESTION 11

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 12

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 13**

Which of the following statements regarding importing streaming data from InfoSphere Streams into Hadoop is TRUE?

- A. InfoSphere Streams can only write to HDFS not read from HDFS
- B. InfoSphere Streams can only write directly to BigInsights, not other Hadoop distributions like Hortonworks or Cloudera
- C. A Streams developer needs to account for the fact that BigInsights may not be able to absorb the incoming streams at the rate InfoSphere Streams is sending them
- D. Adding a Big Data toolkit operator (for writing to Hadoop) to an InfoSphere Streams Processing Language (SPL) application requires that the SPL application be recompiled

Correct Answer: B

Reference: [https://books.google.com.pk/books?id=6BmLCwAAQBAJandpg=PA15andlpg=PA15anddq=InfoSphere+Streams+can+only+write+directly+to+BigInsights,+not+other+Hadoop+distributions+like+Hortonworks+or+Clouderaandsource=bAX5\\_-opwFandsig=ACfU3U2rhwLKLEes37evjJfMzxYMLP7NxQandhl=enandsa=Xandved=2ahUKEwi6gY2soffoAhWlzaQKHREADDkQ6AEwAHoECBcQJg#v=onepageandq=InfoSphere%20Streams%20can%20only%20write%20directly%20to%20BigInsights%2C%20not%20other%20Hadoop%20distributions%20like%20Hortonworks%20or%20Clouderaandf=false](https://books.google.com.pk/books?id=6BmLCwAAQBAJandpg=PA15andlpg=PA15anddq=InfoSphere+Streams+can+only+write+directly+to+BigInsights,+not+other+Hadoop+distributions+like+Hortonworks+or+Clouderaandsource=bAX5_-opwFandsig=ACfU3U2rhwLKLEes37evjJfMzxYMLP7NxQandhl=enandsa=Xandved=2ahUKEwi6gY2soffoAhWlzaQKHREADDkQ6AEwAHoECBcQJg#v=onepageandq=InfoSphere%20Streams%20can%20only%20write%20directly%20to%20BigInsights%2C%20not%20other%20Hadoop%20distributions%20like%20Hortonworks%20or%20Clouderaandf=false)

**QUESTION 14**

Which expression is making proper use of the "LOAD" function in Pig Latin?

- A. A = LOAD data -> name,age,gpa
- B. A = LOAD(\\data\\,[name, age, gpa]);
- C. A = LOAD \\data\\ -> {name:chararray, age:int, gpa:float};
- D. A = LOAD \\data\\ AS (name:chararray, age:int, gpa:float);

Correct Answer: D

Reference: <https://pig.apache.org/docs/r0.17.0/basic.html>

---

**QUESTION 15**

Which command can be used to move data from HDFS into an existing table in a relational database?

- A. sqoop
- B. hadoop fs -put
- C. hadoop fs -copyToLocal
- D. hadoop fs -copyFromLocal



Correct Answer: A

Reference: [https://www.google.com/url? sa=t&drct=j&dq=andesc=sandsource=webandcd=5andcad=rja&duact=8&andved=2ahUKEwjay5\\_\\_o8TeAhUryIUKHQyiDV](https://www.google.com/url?sa=t&drct=j&dq=andesc=sandsource=webandcd=5andcad=rja&duact=8&andved=2ahUKEwjay5__o8TeAhUryIUKHQyiDV)

[https://www.ibm.com/developerworks/community/files/basic/anonymous/api/library/3d16d834-9481-4741-8921-fe4cb5e9c93c/document/fab7d3a7c-2850-4ee1-821b-f6caa1939310/media/hadoop\\_ibm\\_big\\_insights.ppt](https://www.ibm.com/developerworks/community/files/basic/anonymous/api/library/3d16d834-9481-4741-8921-fe4cb5e9c93c/document/fab7d3a7c-2850-4ee1-821b-f6caa1939310/media/hadoop_ibm_big_insights.ppt)  
andusg=AOvVaw1S0xKgU96Ljobe3K3D-1bA

[C2090-101 PDF Dumps](#)

[C2090-101 Practice Test](#)

[C2090-101 Brainsdumps](#)