

# 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.certbus.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

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 2

Which BigInsights tool is able to export data from Big SQL?

- A. Ambari
- B. BigSheets
- C. Flume
- D. Knox

Correct Answer: B

Reference: [https://www.ibm.com/support/knowledgecenter/en/SSCRJT\\_5.0.1/com.ibm.swg.im.bigsql.tut.doc/doc/less\\_bsql\\_exptobigsh.html](https://www.ibm.com/support/knowledgecenter/en/SSCRJT_5.0.1/com.ibm.swg.im.bigsql.tut.doc/doc/less_bsql_exptobigsh.html)

---

### QUESTION 3

What command is used to find the version level of GPFS running on a particular Linux node?

- A. lspp -ver \\*
- B. list ver gpfs\\*
- C. lspp -l gpfs\\*
- D. rpm -qa | grep gpfs

Correct Answer: D

Reference: [https://www.google.com/url?sa=t&andrc=jandq=andesrc=sandsource=webandcd=3andved=2ahUKEwjKhrG7q\\_foAhXiVhUIHSizDyQQFjACegQIAxA&Bandurl=https%3A%2F%2Fwww.ibm.com%2Fdeveloperworks%2Fcommunity%2Fforums%2Fatom%2Fdownload%2FConcepts\\_Planning\\_Installation.pdf%3FnodeId%3Df189c262-a09f-4ed0-991b3e702fafcad2andusg=AOvVaw32QxwHaqCkUPwVyWiS01Cx](https://www.google.com/url?sa=t&andrc=jandq=andesrc=sandsource=webandcd=3andved=2ahUKEwjKhrG7q_foAhXiVhUIHSizDyQQFjACegQIAxA&Bandurl=https%3A%2F%2Fwww.ibm.com%2Fdeveloperworks%2Fcommunity%2Fforums%2Fatom%2Fdownload%2FConcepts_Planning_Installation.pdf%3FnodeId%3Df189c262-a09f-4ed0-991b3e702fafcad2andusg=AOvVaw32QxwHaqCkUPwVyWiS01Cx) (11)

---

#### QUESTION 4

Which is a benefit of row oriented table design?

- A. When writing a new row, if all of the row data is supplied at the same time the entire row can be written with a single disk seek
- B. When columns of a single row are required at the same time, the entire row can be retrieved with a single disk seek regardless of row size
- C. When new values of a column are supplied for all rows at once, that column data can be written efficiently and replace old column data without touching any other columns for the rows
- D. When an aggregate needs to be computed over many rows but only a notably smaller subset of all columns of data, reading that smaller subset of data can be faster than reading all data

Correct Answer: B

Reference: <http://www.ijoart.org/docs/Column-Oriented-Databases-to-Gain-High-Performance-for-DataWarehouse-System.pdf> (7)

---

#### QUESTION 5

When creating a configuration file for a Flume agent, which of the following must be configured?

- A. An interceptor
- B. A database configuration file
- C. A source, a sink, and a channel
- D. All of the above

Correct Answer: C

Flume agents consist of three elements: a source, a channel, and a sink. The channel connects the source to the sink. You must configure each element in the Flume agent. Different source, channel, and sink types have different configurations

Reference: [https://www.ibm.com/support/knowledgecenter/S5CKRH\\_1.0.3/alerts/c\\_fcail\\_modflumeagents.html](https://www.ibm.com/support/knowledgecenter/S5CKRH_1.0.3/alerts/c_fcail_modflumeagents.html)

---

#### QUESTION 6

When running Hive Queries using Apache Spark, which of the following execution engines is invoked?

- A. Tez
- B. Spark
- C. HiveQL

D. MapReduce

Correct Answer: B

When running queries, you can set hive.execution.engine so as to choose the execution engine. To run on Spark, set it to spark. To run on MapReduce, set it to mr.

Reference: [https://www.ibm.com/support/knowledgecenter/en/SSPT3X\\_4.2.5/com.ibm.swg.im.infosphere.biginsights.admin.doc/doc/admin\\_hive\\_on\\_spark.html](https://www.ibm.com/support/knowledgecenter/en/SSPT3X_4.2.5/com.ibm.swg.im.infosphere.biginsights.admin.doc/doc/admin_hive_on_spark.html)

---

### QUESTION 7

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

---

### QUESTION 8

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 9

How are insights derived from Big Match moved to an MDM system?

- A. Extract insights from HBase and load into MDM through an API call
- B. Extract insights from Hive and load into MDM using standard tooling
- C. Extract insights from HDFS and load into MDM by stimulating delta load
- D. Extract insights from HBase and load into MDM using standard MDM batch processing tool

Correct Answer: C

---

#### QUESTION 10

What are the available document formats beside PDF and MS Word when export a redacted document using Optim Review Tool?

- A. TIFF, and CSF
- B. TIFF, and PNG
- C. JPEG, and PNG
- D. Plain Text, and CSV

Correct Answer: A

Reference: [https://www-01.ibm.com/software/info/channel-solution-profiles/Information\\_Graphics\\_RedactIt.html](https://www-01.ibm.com/software/info/channel-solution-profiles/Information_Graphics_RedactIt.html)

---

#### QUESTION 11

Which Big SQL file format could be expected to result in a longer running query?

- A. Text
- B. Avro
- C. Parquet
- D. Sequence\_text

Correct Answer: C

Reference: <https://www.ibm.com/developerworks/data/library/techarticle/dm-1510-parquet-big-sql/index.html>

---

#### QUESTION 12

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

- A. InfoSphere Streams utilized Spark Streaming to interface to Hadoop

- B. The HDFS operators in InfoSphere Streams use Hadoop Java APIs to access HDFS or GPFS
- C. Incoming streams from InfoSphere Streams into Hadoop are not buffered to ensure low latency
- D. Since InfoSphere Streams processes data in Streams tuple format, Hadoop must store the data in this format

Correct Answer: B

Reference: [https://www.ibm.com/support/knowledgecenter/en/SSCRJU\\_4.0.1/com.ibm.streams.toolkits.doc/doc/tk\\$com.ibm.streamsx.hdfs/tk\\$com.ibm.streamsx.hdfs.html](https://www.ibm.com/support/knowledgecenter/en/SSCRJU_4.0.1/com.ibm.streams.toolkits.doc/doc/tk$com.ibm.streamsx.hdfs/tk$com.ibm.streamsx.hdfs.html)

---

### QUESTION 13

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

---

### QUESTION 14

A large Telecom company wants to store data from multiple databases into Hadoop. They plan to do bulk loads of data into Hadoop and run analytical queries. Which data store would be ideal for this scenario?

- A. Hive
- B. HBase
- C. BigSheets
- D. Apache Spark

Correct Answer: B

Reference: <https://developer.ibm.com/recipes/tutorials/big-data-and-hadoop-on-ibm-cloud/>

---

### QUESTION 15

The input to the Reduce function is \_\_\_\_\_.

- A. Sorted by key
- B. Sorted by value
- C. Grouped by mapper
- D. Shuffled into random order

Correct Answer: A

Reference: [https://www.ibm.com/support/knowledgecenter/SSZUMP\\_7.3.0/foundations\\_sym/symphony\\_developer\\_objects.html](https://www.ibm.com/support/knowledgecenter/SSZUMP_7.3.0/foundations_sym/symphony_developer_objects.html)

[C2090-101 Practice Test](#)

[C2090-101 Study Guide](#)

[C2090-101 Braindumps](#)