



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.passapply.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

Consider the following Solr query:

```
curl "http://localhost:8983/solr/gettingstarted/select?wt=json&indent=true&q=foundation"
```

What is the term that is being searched?

- A. indent
- B. json
- C. gettingstarted
- D. foundation

Correct Answer: D

Reference: https://lucene.apache.org/solr/6_0_0/quickstart.html

QUESTION 2

Suppose that you have some log files that you need to load into HBase.

What tool could you use to perform a bulk load of the log files?

- A. Import
- B. Fastload
- C. ImportTsv
- D. None of the above

Correct Answer: C

Reference: https://www.ibm.com/support/knowledgecenter/SSWSR9_11.6.0/com.ibm.swg.im.mdmhs.pmebi.doc/topics/running_automatic.html

QUESTION 3

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 4

What is the primary purpose of Flume in the Hadoop ecosystem?

- A. To stream data from Hadoop
- B. To move static files from the local file system into HDFS
- C. To import data from a relational database or data warehouse into HDFS
- D. To capture log data as it is written to log files and move them into HDFS

Correct Answer: C

Reference: https://blogs.apache.org/flume/entry/flume_ng_architecture

QUESTION 5

Which component of BigInsights is able to mask data items so as restrict viewing of sensitive data?

- A. Flume
- B. HDFS
- C. Oozie
- D. Big SQL

Correct Answer: D

Reference: <https://developer.ibm.com/hadoop/docs/biginsights-ibm-open-platform/biginsights-value-add/big-sql/maximize-hadoop-data-security-ibm-infosphere-biginsights/>

QUESTION 6

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=andesrc=s&source=web&cd=5&cad=rja&uact=8&ved=2ahUKEwjay5__o8TeAhUrylUKHQyiDVYQFjAEegQIBBAC&url=https%3A%2F%2Fwww.ibm.com%2Fdeveloperworks%2Fcommunity%2Ffiles%2Fbasic%2Fanonymous%2Fapi%2Flibrary%2F3d16d834-9481-4741-8921-fe4cb5e9c93c%2Fdocument%2Fab7d3a7c-2850-4ee1-821b-f6caa1939310%2Fmedia%2FHadoop_IBM_Big_Insights.ppt&usq=AOvVaw1S0xKgU96Ljobe3K3D-1bA

QUESTION 7

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 8

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

- A. InfoSphere Streams can both read from and write data to HDFS
- B. The Streams Big Data toolkit operators that interface with HDFS uses Apache Flume to integrate with Hadoop
- C. Streams applications never need to be concerned with making the data schemas consistent with those on Hadoop
- D. Big SQL can be used to preprocess the data as it flows through InfoSphere Streams before the data lands in HDFS

Correct Answer: D

QUESTION 9

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 10

Which of the following is TRUE about storing an Apache Spark object in serialized form?

- A. It is advised to use Java serialization over Kryo serialization
- B. Storing the object in serialized form will lead to faster access times
- C. Storing the object in serialized form will lead to slower access times
- D. All of the above

Correct Answer: B

Reference: <https://spark.apache.org/docs/latest/rdd-programming-guide.html>

QUESTION 11

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 12

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 13

What Redaction feature needs to be selected when manually redacting a form through the Optim Review Tool?

- A. Text Redaction



- B. Image Redaction
- C. Region Redaction
- D. Redact by Information Type

Correct Answer: D

QUESTION 14

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

Correct Answer: CE

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

QUESTION 15

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/>

[C2090-101 PDF Dumps](#)

[C2090-101 VCE Dumps](#)

[C2090-101 Braindumps](#)