

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.leads4pass.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

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 2

Which of the following are capabilities of the Apache Spark project?

- A. Large scale machine learning
- B. Large scale graph processing
- C. Live data stream processing
- D. All of the above

Correct Answer: B

Reference: <https://spark.apache.org/>

QUESTION 3

The number of partitions created by DynamicPartitions in Hive can be controlled by which of the following?

- A. hive.exec.max.dynamic.partitions.pernode
- B. hive.exec.max.dynamic.partitions
- C. hive.exec.max.created.files
- D. All of the above

Correct Answer: A

Reference: <https://resources.zaloni.com/blog/partitioning-in-hive>

QUESTION 4

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

QUESTION 5

Which one of the following statements is TRUE?

- A. Spark SQL does not support HiveQL
- B. Spark SQL does not support ANSI SQL
- C. To use Spark with Hive, HiveQL queries have to be rewritten in Scala
- D. Spark SQL allows relational queries expressed in SQL, HiveQL, or Scala

Correct Answer: D

Reference: <https://spark.apache.org/docs/1.2.1/sql-programming-guide.html> (overview)

QUESTION 6

The GPFS implementation of Data Management API is compliant to which Open Group storage management Standard?

- A. XSH
- B. XBD
- C. XDMS
- D. X / Open

Correct Answer: C

Reference: https://www.ibm.com/support/knowledgecenter/en/SSFKCN_4.1.0/com.ibm.cluster.gpfs.v4r1.gpfs400.doc/bl1dmp_intro.htm

QUESTION 7

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

QUESTION 8

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

Correct Answer: B

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

QUESTION 9

Which of the following statement is TRUE with BigSheets?

- A. You can create any type of sheet from a parent workbook
- B. You must create a child workbook in order to create a chart
- C. You can delete a parent workbook without deleting the child workbooks
- D. You must run the workbook on the data to get the full results of the analysis

Correct Answer: A

Reference: https://www.ibm.com/support/knowledgecenter/SSPT3X_4.1.0/com.ibm.swg.im.infosphere.biginsights.analyze.doc/doc/bigsheets_con_workbooks.html

QUESTION 10

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 11

For what purpose SPSS models are embedded within InfoSphere Streams application?

- A. To provide high availability
- B. To score streaming data using existing models
- C. To create new models based on streaming data
- D. To ingest and parse binary and other complex data types

Correct Answer: B

Reference: <https://www.ibm.com/developerworks/data/tutorials/dm-1109spssscoringinfospherestreams1/dm-1109spssscoringinfospherestreams1-pdf.pdf>

QUESTION 12

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

Correct Answer: DE

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 13

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 built-in 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`

Correct Answer: C

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

QUESTION 14

What is a standard method of monitoring sensitive data in Hadoop?

A. Use any database monitoring tool

B. Use HBASE GET or `hadoop fs -cat` commands

C. Use STAP on each Hadoop Node

D. All of the above

Correct Answer: C

Reference: <https://www.ibm.com/developerworks/data/library/techarticle/dm-1210bigdatasecurity/index.html>

QUESTION 15

Which of the following statements regarding Big SQL is TRUE?

A. Big SQL doesn't support stored procedures

B. Big SQL can be deployed on a subset of data nodes in the BigInsights cluster

C. Big SQL provides a SQL-on-Hadoop environment based on map reduce

D. Only tables created or loaded via Big SQL can be accessed via Big SQL

Correct Answer: B

Reference: https://books.google.com.pk/books?id=t13nCQAAQBAJandpg=PA3andlpg=PA3anddq=Big+SQL+can+be+deployed+on+a+subset+of+data+nodes+in+the+BigInsights+clusterandsource=blandots=RBbad0Xkelandsig=pMgmgDNLGUrkvOSXoVBj64xTMgkandhl=enandsa=Xandredir_esc=y#v=onepageandq=Big%20SQL%20can%20be%20deployed%20on%20a%20subset%20of%20data%20nodes%20in%20the%20BigInsights%20clusterandf=false

[C2090-101 PDF Dumps](#)

[C2090-101 Practice Test](#)

[C2090-101 Braindumps](#)