

# QUESTIONS & ANSWERS

Kill your exam at first Attempt



C2090-101 Dumps  
C2090-101 Braindumps  
C2090-101 Real Questions  
C2090-101 Practice Test  
C2090-101 dumps free



**IBM**

# C2090-101

*IBM Big Data Engineer*

<http://killexams.com/pass4sure/exam-detail/C2090-101>



Question: 95 Section 1

Which of the following is NOT TRUE regarding InfoSphere Streams composite operators that are annotated as consistent regions?

- A. Disconnected subgraphs within the composite participate in the consistent region
- B. The composite operator and the reachability graph of the composite operator participate in a single consistent region
- C. Merging composite operators into a single consistent region can yield unexpected results and is therefore not recommended
- D. Operators outside of the composite may participate in the consistent region if they are within the composite's reachability graph

**Answer: C**

Reference:

[https://www.ibm.com/support/knowledgecenter/SSCRJU\\_4.0.0/com.ibm.streams.dev.doc/doc/consistentcompositeoperators.html](https://www.ibm.com/support/knowledgecenter/SSCRJU_4.0.0/com.ibm.streams.dev.doc/doc/consistentcompositeoperators.html)

Question: 96 Section 1

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

**Answer: C**

Reference:

[https://www.ibm.com/support/knowledgecenter/SSWSR9\\_11.6.0/com.ibm.swg.im.mdmhs.pmebi.doc/topics/running\\_automatic.html](https://www.ibm.com/support/knowledgecenter/SSWSR9_11.6.0/com.ibm.swg.im.mdmhs.pmebi.doc/topics/running_automatic.html)

Question: 97 Section 1

What should you use to run Big SQL queries that require dynamic variables or scheduling?

- A. JSqsh
- B. Eclipse
- C. BigInsights Console
- D. BigInsights Applications

**Answer: D**

Reference:

[https://www.ibm.com/support/knowledgecenter/SSPT3X\\_3.0.0/com.ibm.swg.im.infosphere.biginsights.analyze.doc/doc/bigsql\\_run\\_queries.html](https://www.ibm.com/support/knowledgecenter/SSPT3X_3.0.0/com.ibm.swg.im.infosphere.biginsights.analyze.doc/doc/bigsql_run_queries.html)

Question: 98 Section 1

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`

**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](https://www.ibm.com/support/knowledgecenter/SSPT3X_3.0.0/com.ibm.swg.im.infosphere.biginsights.import.doc/doc/data_warehouse_sqoop.html)

Question: 99 Section 1

You have social media data in a plain text file. You want to figure out the number of times the word "Watson" is mentioned in it.

Which of the following steps will correctly allow you to visualize that data?

- A. Create a master workbook, add a filter sheet, run the workbook, add chart
- B. Create a master workbook, create a child workbook, add a filter sheet, add chart
- C. Create a master workbook, edit the workbook, add a filter sheet, run the workbook, add chart
- D. Create a master workbook, create a child workbook, add a filter sheet, run the workbook, add chart

**Answer: C**

Question: 100 Section 1

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

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

**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: 101 Section 1

A Resilient Distributed Dataset is a good solution for which of the following?

- A. A search
- B. An SQL based data transformation
- C. A high throughput write oriented use cases
- D. An interactive query and iterative algorithms

**Answer: D**

Reference:  
<https://www.usenix.org/system/files/conference/nsdi12/nsdi12-final138.pdf>

Question: 102 Section 1

A large retailer is planning to create a data cloud using BigInsights. They would like to run queries using data that is stored in BigInsights as well as other databases like DB2 and Oracle.

Which one of the following would enable them to write the queries without having to move data into one common data store?

- A. Hive
- B. HBase
- C. Big SQL
- D. BigSheets

**Answer: C**

Reference:  
<https://www.ibmdatahub.com/blog/hadoop-meets-sql>

Question: 103 Section 1

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

**Answer: B**

Question: 104 Section 1

How many Job Trackers can be found in a MapReduce v1 cluster?

- A. One per cluster
- B. One per data node
- C. One for each Mapper
- D. One for each Reducer

**Answer: A**

Reference:  
<https://www.cs.colorado.edu/~kena/classes/5448/s11/presentations/hadoop.pdf>

Question: 105 Section 1

Which of the following is NOT TRUE regarding the Job Control Plane operator?

- A. User-defined checkpoint logic can be coded in the logic section
- B. It must be added to each application that includes a consistent region
- C. It controls the notifications to set of operators that are included in a consistent region
- D. It has no input or output ports and appears as a stand-alone operator in an application graph

**Answer: A**

Reference:  
[https://www.ibm.com/support/knowledgecenter/SSCRJU\\_4.3.0/com.ibm.streams.dev.doc/doc/jobcontrolplane.html](https://www.ibm.com/support/knowledgecenter/SSCRJU_4.3.0/com.ibm.streams.dev.doc/doc/jobcontrolplane.html)

Question: 106 Section 1

Which one of the following file formats is optimal for querying tables with many columns and performing aggregation operations such as SUM() and AVG()?

- A. Text
- B. Avro
- C. JSON
- D. PARQUET

Answer: C



For More exams visit <https://killexams.com/vendors-exam-list>



*Kill your exam at First Attempt....Guaranteed!*