



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.geekcert.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

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

Correct Answer: C

QUESTION 2

Which ONE of the following statements regarding Sqoop is TRUE?

- A. By default, data is compressed with Sqoop
- B. Sqoop can only read committed transactions from a source database, not uncommitted ones
- C. When performing parallel imports, Sqoop always uses the primary key column in a table as the splitting column
- D. When performing parallel imports, each degree of parallelism corresponds to a concurrent database connection

Correct Answer: B

By default, Sqoop uses the read committed transaction isolation in the mappers to import data. This may not be the ideal in all ETL workflows and it may be desired to reduce the isolation guarantees. The `--relaxedisolation` option can be used to instruct Sqoop to use read uncommitted isolation level.

Reference: <https://sqoop.apache.org/docs/1.4.6/SqoopUserGuide.html>

QUESTION 3

A Resilient Distributed Dataset supports which of the following?

- A. Creating a new dataset from an old one
- B. Returning a computed value to the driver program
- C. Both "Creating a new dataset from an old one" and "Returning a computed value to the driver program"
- D. Neither "Creating a new dataset from an old one" nor "Returning a computed value to the driver program"

Correct Answer: C



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

QUESTION 4

What is a method for loading RDBMS data into an HBase table?

- A. HDFS LOAD
- B. SQOOP IMPORT
- C. LOAD HADOOP USING
- D. Hadoop jar hbase-VERSION.jar importtsv

Correct Answer: B

QUESTION 5

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/>

[Latest C2090-101 Dumps](#)

[C2090-101 Practice Test](#)

[C2090-101 Exam Questions](#)