



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

Which Big SQL statement can be used to store a single row into an HBase table?

- A. LOAD HIVE DATA
- B. INSERT INTO HBASE
- C. CREATE HBASE TABLE
- D. LOAD USING ... INTO HBASE TABLE

Correct Answer: B

Reference: https://www.ibm.com/support/knowledgecenter/en/SSPT3X_3.0.0/com.ibm.swg.im.infosphere.biginsights.bigsql.doc/doc/bsql_insert.html

QUESTION 2

Which of the following is not a capability of Pig?

- A. Low-latency queries
- B. Schemas are optional
- C. Nested relational data model
- D. A high level abstraction on top of MapReduce

Correct Answer: A

Reference: <http://hadooptutorial.info/apache-pig-overview/>

QUESTION 3

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

Correct Answer: A

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



QUESTION 4

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 5

A large bank was planning to offload existing data from a data warehouse into Hadoop and use SQL queries to access historical data. Which one of the following statements is true for using HiveQL?

- A. It supports four logical operators in query predicates: IN, NOT IN, EXISTS, and NOT EXISTS
- B. It does not support nested sub-queries
- C. Hive supports all ANSI SQL 2011 syntax
- D. All of the above

Correct Answer: A

Reference: https://docs.hortonworks.com/HDPDocuments/HDP2/HDP-2.6.5/bk_data-access/content/hive013-feature-subqueries-in-where-clauses.html

[C2090-101 PDF Dumps](#)

[C2090-101 Study Guide](#)

[C2090-101 Braindumps](#)