

APACHE-HADOOP-DEVELOPER^{Q&As}

Hadoop 2.0 Certification exam for Pig and Hive Developer

Pass Hortonworks APACHE-HADOOP-DEVELOPER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass2lead.com/apache-hadoop-developer.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by
Hortonworks Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



QUESTION 1

All keys used for intermediate output from mappers must:

- A. Implement a splittable compression algorithm.
- B. Be a subclass of FileInputFormat.
- C. Implement WritableComparable.
- D. Override isSplittable.
- E. Implement a comparator for speedy sorting.

Correct Answer: C

Explanation: The MapReduce framework operates exclusively on pairs, that is, the framework views the input to the job as a set of pairs and produces a set of pairs as the output of the job, conceivably of different types.

The key and value classes have to be serializable by the framework and hence need to implement the Writable interface. Additionally, the key classes have to implement the WritableComparable interface to facilitate sorting by the framework.

Reference: MapReduce Tutorial

QUESTION 2

Given the following Hive commands:

```
CREATE TABLE mytable (name chararray, age int) ROW FORMAT
DELIMITED FIELDS TERMINATED BY ',' STORED AS
TEXTFILE;
LOAD DATA INPATH '/home/user/mydata.txt' INTO TABLE mytable;
```

Which one of the following statements is true?

- A. The file mydata.txt is copied to a subfolder of /apps/hive/warehouse
- B. The file mydata.txt is moved to a subfolder of /apps/hive/warehouse
- C. The file mydata.txt is copied into Hive's underlying relational database.
- D. The file mydata.txt does not move from its current location in HDFS

Correct Answer: A

QUESTION 3

Which two of the following are true about this trivial Pig program? (choose Two)

```
$ pig
grunt> ABC = LOAD 'myfile';
grunt> DUMP ABC;
```

- A. The contents of myfile appear on stdout
- B. Pig assumes the contents of myfile are comma delimited
- C. ABC has a schema associated with it
- D. myfile is read from the user's home directory in HDFS

Correct Answer: AD

QUESTION 4

In a MapReduce job, the reducer receives all values associated with same key. Which statement best describes the ordering of these values?

- A. The values are in sorted order.
- B. The values are arbitrarily ordered, and the ordering may vary from run to run of the same MapReduce job.
- C. The values are arbitrary ordered, but multiple runs of the same MapReduce job will always have the same ordering.
- D. Since the values come from mapper outputs, the reducers will receive contiguous sections of sorted values.

Correct Answer: B

Note:

*

Input to the Reducer is the sorted output of the mappers.

*

The framework calls the application's Reduce function once for each unique key in the sorted order.

*

Example:

For the given sample input the first map emits:

The second map emits:

QUESTION 5

To use a java user-defined function (UDF) with Pig what must you do?

- A. Define an alias to shorten the function name
- B. Pass arguments to the constructor of UDFs implementation class
- C. Register the JAR file containing the UDF
- D. Put the JAR file into the user's home folder in HDFS

Correct Answer: C

[Latest APACHE-HADOOP-DEVELOPER Dumps](#)

[APACHE-HADOOP-DEVELOPER VCE Dumps](#)

[APACHE-HADOOP-DEVELOPER Study Guide](#)