



HADOOP-PR000007^{Q&As}

Hortonworks Certified Apache Hadoop 2.0 Developer (Pig and Hive Developer)

Pass Hortonworks HADOOP-PR000007 Exam with 100% Guarantee

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

<https://www.pass4itsure.com/hadoop-pr000007.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**

Which one of the following statements is true regarding a MapReduce job?

- A. The job's Partitioner shuffles and sorts all (key.value) pairs and sends the output to all reducers
- B. The default Hash Partitioner sends key value pairs with the same key to the same Reducer
- C. The reduce method is invoked once for each unique value
- D. The Mapper must sort its output of (key.value) pairs in descending order based on value

Correct Answer: A

QUESTION 2

Review the following andapos;dataandapos; file and Pig code.

```
M,38,95111
```

```
M,62,95102
```

```
A = LOAD andapos;dataandapos; USING PigStorage(andapos;,andapos;)  
AS (gender:chararray, age:int, zip:chararray);
```

```
D = GROUP A BY gender;  
DUMP D;
```

Which one of the following statements is true?

- A. The Output Of the DUMP D command IS (M,{{(M,62.95102),(M,38,95111}})
- B. The output of the dump d command is (M, {{(38,95in),(62,95i02}})
- C. The code executes successfully but there is not output because the D relation is empty
- D. The code does not execute successfully because D is not a valid relation

Correct Answer: A

QUESTION 3

You need to perform statistical analysis in your MapReduce job and would like to call methods in the Apache Commons Math library, which is distributed as a 1.3 megabyte Java archive (JAR) file. Which is the best way to make this library available to your MapReducer job at runtime?

- A. Have your system administrator copy the JAR to all nodes in the cluster and set its location in the HADOOP_CLASSPATH environment variable before you submit your job.



- B. Have your system administrator place the JAR file on a Web server accessible to all cluster nodes and then set the HTTP_JAR_URL environment variable to its location.
- C. When submitting the job on the command line, specify the ?ibjars option followed by the JAR file path.
- D. Package your code and the Apache Commons Math library into a zip file named JobJar.zip

Correct Answer: C

Explanation: The usage of the jar command is like this,

Usage: `hadoop jar [mainClass] args...`

If you want the commons-math3.jar to be available for all the tasks you can do any one of these

1. Copy the jar file in \$HADOOP_HOME/lib dir

2.

or

Use the generic option `-libjars`.

QUESTION 4

Which one of the following statements describes a Hive user-defined aggregate function?

- A. Operates on multiple input rows and creates a single row as output
- B. Operates on a single input row and produces a single row as output
- C. Operates on a single input row and produces a table as output
- D. Operates on multiple input rows and produces a table as output

Correct Answer: A

QUESTION 5

Given the following Hive command:

```
INSERT OVERWRITE TABLE mytable SELECT * FROM myothertable;
```

Which one of the following statements is true?

- A. The contents of myothertable are appended to mytable
- B. Any existing data in mytable will be overwritten
- C. A new table named mytable is created, and the contents of myothertable are copied into mytable
-



D. The statement is not a valid Hive command

Correct Answer: B

[HADOOP-PR000007 VCE Dumps](#)

[HADOOP-PR000007 Practice Test](#)

[HADOOP-PR000007 Braindumps](#)