



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.passapply.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

In a MapReduce job, you want each of your input files processed by a single map task. How do you configure a MapReduce job so that a single map task processes each input file regardless of how many blocks the input file occupies?

- A. Increase the parameter that controls minimum split size in the job configuration.
- B. Write a custom MapRunner that iterates over all key-value pairs in the entire file.
- C. Set the number of mappers equal to the number of input files you want to process.
- D. Write a custom FileInputFormat and override the method isSplittable to always return false.

Correct Answer: D

Explanation: FileInputFormat is the base class for all file-based InputFormats. This provides a generic implementation of getSplits(JobContext). Subclasses of FileInputFormat can also override the isSplittable (JobContext, Path) method to ensure input-files are not split-up and are processed as a whole by Mappers.

Reference: org.apache.hadoop.mapreduce.lib.input, Class FileInputFormat

QUESTION 2

Which one of the following statements is FALSE regarding the communication between DataNodes and a federation of NameNodes in Hadoop 2.0?

- A. Each DataNode receives commands from one designated master NameNode.
- B. DataNodes send periodic heartbeats to all the NameNodes.
- C. Each DataNode registers with all the NameNodes.
- D. DataNodes send periodic block reports to all the NameNodes.

Correct Answer: A

QUESTION 3

Examine the following Hive statements:

```
CREATE TABLE x (name STRING, age INT, zip INT, salary DOUBLE)
ROW FORMAT DELIMITED FIELDS TERMINATED BY
',';
LOAD DATA INPATH 'input/File1' OVERWRITE INTO TABLE x;
SELECT * FROM x SORT BY age;
```

Assuming the statements above execute successfully, which one of the following statements is true?



- A. Each reducer generates a file sorted by age
- B. The SORT BY command causes only one reducer to be used
- C. The output of each reducer is only the age column
- D. The output is guaranteed to be a single file with all the data sorted by age

Correct Answer: A

QUESTION 4

Consider the following two relations, A and B.

```
A = LOAD 'data1' AS (a1:int,a2:chararray);
DUMP A;
(1,apple)
(3,orange)
(4,peach)
(2,cherry)

B = LOAD 'data2' AS (b1:chararray,b2:int);
DUMP B;
(Jim,2)
(Brian,4)
(Kim,0)
(Terry,3)
(Chris,2)
```

A Pig JOIN statement that combined relations A by its first field and B by its second field would produce what output?

- A. 2 Jim Chris 2 3 Terry 3 4 Brian 4
- B. 2 cherry 2 cherry 3 orange 4 peach
- C. 2 cherry Jim, Chris 3 orange Terry
4 peach Brian
- D. 2 cherry Jim 2 2 cherry Chris 2 3 orange Terry 3 4 peach Brian 4

Correct Answer: D

QUESTION 5

In Hadoop 2.0, which one of the following statements is true about a standby NameNode? The Standby NameNode:

- A. Communicates directly with the active NameNode to maintain the state of the active NameNode.
- B. Receives the same block reports as the active NameNode.
- C. Runs on the same machine and shares the memory of the active NameNode.



D. Processes all client requests and block reports from the appropriate DataNodes.

Correct Answer: B

[Latest HADOOP-PR000007 Dumps](#)

[HADOOP-PR000007 VCE Dumps](#)

[HADOOP-PR000007 Study Guide](#)