

# HADOOP-PR000007<sup>Q&As</sup>

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.pass2lead.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

You wrote a map function that throws a runtime exception when it encounters a control character in input data. The input supplied to your mapper contains twelve such characters total, spread across five file splits. The first four file splits each have two control characters and the last split has four control characters.

Identify the number of failed task attempts you can expect when you run the job with `mapred.max.map.attempts` set to 4:

- A. You will have forty-eight failed task attempts
- B. You will have seventeen failed task attempts
- C. You will have five failed task attempts
- D. You will have twelve failed task attempts
- E. You will have twenty failed task attempts

Correct Answer: E

Explanation: There will be four failed task attempts for each of the five file splits.

Note:

When the jobtracker is notified of a task attempt that has failed (by the tasktracker's heartbeat call), it will reschedule execution of the task. The jobtracker will try to avoid rescheduling the task on a tasktracker where it has previously failed. Furthermore, if a task fails four times (or more), it will not be retried further. This value is configurable: the maximum number of attempts to run a task is controlled by the `mapred.map.max.attempts` property for map tasks and `mapred.reduce.max.attempts` for reduce tasks. By default, if any task fails four times (or whatever the maximum number of attempts is configured to), the whole job fails.

---

### QUESTION 2

Which one of the following Hive commands uses an HCatalog table named x?

- A. `SELECT * FROM x;`
- B. `SELECT x.-FROM org.apache.hcatalog.hive.HCatLoader(\\x\\);`
- C. `SELECT * FROM org.apache.hcatalog.hive.HCatLoader(\\x\\);`
- D. Hive commands cannot reference an HCatalog table

Correct Answer: C

---

### QUESTION 3

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

#### QUESTION 4

Given a directory of files with the following structure: line number, tab character, string: Example:

1abialkjkfjkaosdfjksdlkjhweroij 2kadfjhuwqounahagtnbvaswslmnbfgy 3kjfteiomndscxeqalkzhtopedkfsikj You want to send each line as one record to your Mapper. Which InputFormat should you use to complete

the line: `conf.setInputFormat (____.class) ; ?`

- A. SequenceFileAsTextInputFormat
- B. SequenceFileInputFormat
- C. KeyValueFileInputFormat
- D. BDBInputFormat

Correct Answer: C

<http://stackoverflow.com/questions/9721754/how-to-parse-customwritable-from-text-in-hadoop>

#### QUESTION 5

What is the term for the process of moving map outputs to the reducers?

- A. Reducing
- B. Combining
- C. Partitioning
- D. Shuffling and sorting

Correct Answer: D

[Latest HADOOP-PR000007 Dumps](#)

[HADOOP-PR000007 VCE Dumps](#)

[HADOOP-PR000007 Exam Questions](#)