

TorrentExam

Input your exam code ...

The passing rate of our valid exam braindumps for most certifications is high up to 99%. A small PDF dumps free is ready to download for new customers to tell if our exam dumps are suitable for their real exam.

All Products | Contact now

Why Choose Us



QUALITY AND VALUE

RealExamFree Practice Exams are written to the highest standards of technical accuracy, using only certified subject matter experts and published authors for development - no all dumps.



TESTED AND APPROVED

We are committed to the process of vendor and third party approvals. We believe professionals and executives alike deserve the confidence of quality coverage these authorizations provide.



EASY TO PASS

If you prepare for the exams using our RealExamFree testing engine, It is easy to succeed for all certifications in the first attempt. You don't have to deal with all dumps or any free torrent / rapidshare all stuff.



TRY BEFORE BUY

RealExamFree offers free demo of each product. You can check out the interface, question quality and usability of our practice exams before you decide to buy.

Customer Reviews



I wasted a lot of money and failed twice. Thanks to HPE0-J78 exam collection I pass now.

Noel



realexamfree is a reliable company. I pass exam at first shot. Many thanks

Julie



Pass BIMF.EN successfully. Really good dumps. It saves me a lot of time. Wonderful!

Ahern



online test engine is very useful for me, because i could practice the C-TERP10-67 question dumps in my phone when i was waiting or on the bus even without internet, i could make the most of my time. Last week, i passed the C-TERP10-67. so i want to share the realexamfree with you guys. hope you will get a good result in test.

Carl

<http://www.torrentexam.com>

Best Exam Bootcamp & Excellent VCE Torrent & Satisfying Dumps Torrent

Exam : CCD-333

Title : Cloudera Certified Developer for
Apache Hadoop

Vendors : Cloudera

Version : DEMO

NO.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.

Answer: B

NO.2 You need to import a portion of a relational database every day as files to HDFS, and generate Java classes to interact with your imported data. Which of the following tools should you use to accomplish this?

- A. Pig
- B. Hue
- C. Hive
- D. Flume
- E. Sqoop
- F. Oozie
- G. fuse-dfs

Answer: C,E

NO.3 Which of the following statements most accurately describes the relationship between MapReduce and Pig?

- A. Pig provides additional capabilities that allow certain types of data manipulation not possible with MapReduce.
- B. Pig provides no additional capabilities to MapReduce. Pig programs are executed as MapReduce jobs via the Pig interpreter.
- C. Pig programs rely on MapReduce but are extensible, allowing developers to do special-purpose processing not provided by MapReduce.

D. Pig provides the additional capability of allowing you to control the flow of multiple MapReduce jobs.

Answer: D

NO.4 What is a SequenceFile?

A. A SequenceFile contains a binary encoding of an arbitrary number of homogeneous writable objects.

B. A SequenceFile contains a binary encoding of an arbitrary number of heterogeneous writable objects.

C. A SequenceFile contains a binary encoding of an arbitrary number of WritableComparable objects, in sorted order.

D. A SequenceFile contains a binary encoding of an arbitrary number key-value pairs. Each key must be the same type. Each value must be the same type.

Answer: D

NO.5 Which of the following utilities allows you to create and run MapReduce jobs with any executable or script as the mapper and/or the reducer?

A. Oozie

B. Sqoop

C. Flume

D. Hadoop Streaming

Answer: D

NO.6 You need a distributed, scalable, data Store that allows you random, realtime read/write access to hundreds of terabytes of data. Which of the following would you use?

A. Hue

B. Pig

C. Hive

D. Oozie

E. HBase

F. Flume

G. Sqoop

Answer: E

NO.7 You have an employee who is a Data Analyst and is very comfortable with SQL. He would like to run ad-hoc analysis on data in your HDFS cluster. Which of the following is a data warehousing software built on top of Apache Hadoop that defines a simple SQL-like query language well-suited for this kind of user?

- A. Pig
- B. Hue
- C. Hive
- D. Sqoop
- E. Oozie
- F. Flume
- G. Hadoop Streaming

Answer: C

NO.8 Which of the following best describes the workings of TextInputFormat?

- A. Input file splits may cross line breaks. A line that crosses file splits is ignored.
- B. The input file is split exactly at the line breaks, so each Record Reader will read a series of complete lines.
- C. Input file splits may cross line breaks. A line that crosses file splits is read by the RecordReaders of both splits containing the brokenline.
- D. Input file splits may cross line breaks. A line that crosses file splits is read by the RecordReader of the split that contains the end of the brokenline.
- E. Input file splits may cross line breaks. A line that crosses file splits is read by the RecordReader of the split that contains the beginning of the broken line.

Answer: D

NO.9 Given a directory of files with the following structure: line number, tab character, string:

Example:

```
abialkjfjkaoasdfjksdlkjhweroij
kadf jhuwqounahagtnbvaswslmnbfgy
kjfteiomndscxeqalkzhtopedkfslkj
```

You want to send each line as one record to your Mapper. Which InputFormat would you use to complete

the line: setInputFormat (_____.class);

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

Answer: C

NO.10 Workflows expressed in Oozie can contain:

- A. Iterative repetition of MapReduce jobs until a desired answer or state is reached.
- B. Sequences of MapReduce and Pig jobs. These are limited to linear sequences of actions with exception handlers but no forks.
- C. Sequences of MapReduce jobs only; no Pig or Hive tasks or jobs. These MapReduce sequences can be combined with forks and path joins.
- D. Sequences of MapReduce and Pig. These sequences can be combined with other actions including forks, decision points, and path joins.

Answer: D