

TestkingPDF



Try before you buy

Download a free sample of any of our exam questions and answers

 Download Demo

- ✓ 24/7 customer support, Secure shopping site
- ✓ Free One year updates to match real exam scenarios
- ✓ If you failed your exam after buying our products we will refund the full amount back to you.



Quality and Value

TestkingPDF 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 TestkingPDF 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

TestkingPDF 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.

<http://www.testkingpdf.com>

Valid test online & stable pass king & latest PDF dumps

Exam : **CCD-470**

Title : Cloudera Certified
Developer for Apache
Hadoop CDH4 Upgrade
Exam (CCDH)

Vendors : Cloudera

Version : DEMO

NO.1 What data does a Reducer reduce method process?

- A. All the data in a single input file.
- B. All data produced by a single mapper.
- C. All data for a given key, regardless of which mapper(s) produced it.
- D. All data for a given value, regardless of which mapper(s) produced it.

Answer: C

NO.2 What happens in a MapReduce job when you set the number of reducers to zero?

- A. No reducer executes, but the mappers generate no output.
- B. No reducer executes, and the output of each mapper is written to a separate file in HDFS.
- C. No reducer executes, but the outputs of all the mappers are gathered together and written to a single file in HDFS.
- D. Setting the number of reducers to zero is invalid, and an exception is thrown.

Answer: B

NO.3 Identify the tool best suited to import a portion of a relational database every day as files into HDFS, and generate Java classes to interact with that imported data?

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

Answer: F

NO.4 You want to populate an associative array in order to perform a map-side join. You've decided to put this information in a text file, place that file into the DistributedCache and read it in your Mapper before any records are processed.

Identify which method in the Mapper you should use to implement code for reading the file and populating the associative array?

- A. combine
- B. map
- C. init
- D. configure

Answer: D

NO.5 Which describes how a client reads a file from HDFS?

- A. The client queries the NameNode for the block location(s). The NameNode returns the block location(s) to the client. The client reads the data directory off the DataNode(s).
- B. The client queries all DataNodes in parallel. The DataNode that contains the requested data responds directly to the client. The client reads the data directly off the DataNode.
- C. The client contacts the NameNode for the block location(s). The NameNode then queries the

DataNodes for block locations. The DataNodes respond to the NameNode, and the NameNode redirects the client to the DataNode that holds the requested data block(s). The client then reads the data directly off the DataNode.

D. The client contacts the NameNode for the block location(s). The NameNode contacts the DataNode that holds the requested data block. Data is transferred from the DataNode to the NameNode, and then from the NameNode to the client.

Answer: C

NO.6 You are running a job that will process a single InputSplit on a cluster which has no other jobs currently running. Each node has an equal number of open Map slots. On which node will Hadoop first attempt to run the Map task?

- A. The node with the most memory
- B. The node with the lowest system load
- C. The node on which this InputSplit is stored
- D. The node with the most free local disk space

Answer: C

NO.7 Assuming default settings, which best describes the order of data provided to a reducer's reduce method:

- A. The keys given to a reducer aren't in a predictable order, but the values associated with those keys always are.
- B. Both the keys and values passed to a reducer always appear in sorted order.
- C. Neither keys nor values are in any predictable order.
- D. The keys given to a reducer are in sorted order but the values associated with each key are in no predictable order

Answer: D