



# CCA-505<sup>Q&As</sup>

Cloudera Certified Administrator for Apache Hadoop (CCAH) CDH5  
Upgrade Exam

## Pass Cloudera CCA-505 Exam with 100% Guarantee

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

<https://www.passapply.com/cca-505.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by Cloudera  
Official Exam Center

- ⚙ **Instant Download** After Purchase
- ⚙ **100% Money Back** Guarantee
- ⚙ **365 Days** Free Update
- ⚙ **800,000+** Satisfied Customers





### QUESTION 1

You observe that the number of spilled records from Map tasks far exceeds the number of map output records. Your child heap size is 1GB and your `io.sort.mb` value is set to 100 MB. How would you tune your `io.sort.mb` value to achieve maximum memory to disk I/O ratio?

- A. Decrease the `io.sort.mb` value to 0
- B. Increase the `io.sort.mb` to 1GB
- C. For 1GB child heap size an `io.sort.mb` of 128 MB will always maximize memory to disk I/O
- D. Tune the `io.sort.mb` value until you observe that the number of spilled records equals (or is as close to equals) the number of map output records

Correct Answer: D

---

### QUESTION 2

A slave node in your cluster has four 2TB hard drives installed (4 x 2TB). The DataNode is configured to store HDFS blocks on the disks. You set the value of the `dfs.datanode.du.reserved` parameter to 100GB. How does this alter HDFS block storage?

- A. A maximum of 100 GB on each hard drive may be used to store HDFS blocks
- B. All hard drives may be used to store HDFS blocks as long as atleast 100 GB in total is available on the node
- C. 100 GB on each hard drive may not be used to store HDFS blocks
- D. 25 GB on each hard drive may not be used to store HDFS blocks

Correct Answer: B

---

### QUESTION 3

You have installed a cluster running HDFS and MapReduce version 2 (MRv2) on YARN. You have no `dfs.hosts` entry()ies in your `hdfs-alte.xml` configuration file. You configure a new worker node by setting `fs.default.name` in its configuration files to point to the NameNode on your cluster, and you start the DataNode daemon on that worker node.

What do you have to do on the cluster to allow the worker node to join, and start storing HDFS blocks?

- A. Nothing; the worker node will automatically join the cluster when the DataNode daemon is started.
- B. Without creating a `dfs.hosts` file or making any entries, run the command `hadoop dfsadmin refreshHadoop` on the NameNode
- C. Create a `dfs.hosts` file on the NameNode, add the worker node's name to it, then issue the command `hadoop dfsadmin refreshNodes` on the NameNode
- D. Restart the NameNode



Correct Answer: B

---

#### QUESTION 4

You want to understand more about how users browse your public website. For example, you want to know which pages they visit prior to placing an order. You have a server farm of 200 web servers hosting your website. Which is the most efficient process to gather these web server logs into your Hadoop cluster for analysis?

- A. Sample the web server logs from web servers and copy them into HDFS using curl
- B. Ingest the server web logs into HDFS using Flume
- C. Import all users clicks from your OLTP databases into Hadoop using Sqoop
- D. Write a MapReduce job with the web servers from mappers and the Hadoop cluster nodes reducers
- E. Channel these clickstream into Hadoop using Hadoop Streaming

Correct Answer: AB

---

#### QUESTION 5

Your cluster implements HDFS High Availability (HA). Your two NameNodes are named nn01 and nn02. What occurs when you execute the command: `hdfs haadmin failover nn01 nn02`

- A. nn02 becomes the standby NameNode and nn01 becomes the active NameNode
- B. nn02 is fenced, and nn01 becomes the active NameNode
- C. nn01 becomes the standby NameNode and nn02 becomes the active NameNode
- D. nn01 is fenced, and nn02 becomes the active NameNode

Correct Answer: D

[Latest CCA-505 Dumps](#)

[CCA-505 VCE Dumps](#)

[CCA-505 Braindumps](#)