

# CCA-500<sup>Q&As</sup>

Cloudera Certified Administrator for Apache Hadoop (CCA-H)

## Pass Cloudera CCA-500 Exam with 100% Guarantee

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

<https://www.pass2lead.com/cca-500.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 are planning a Hadoop cluster and considering implementing 10 Gigabit Ethernet as the network fabric. Which workloads benefit the most from faster network fabric?

- A. When your workload generates a large amount of output data, significantly larger than the amount of intermediate data
- B. When your workload consumes a large amount of input data, relative to the entire capacity of HDFS
- C. When your workload consists of processor-intensive tasks
- D. When your workload generates a large amount of intermediate data, on the order of the input data itself

Correct Answer: A

---

### QUESTION 2

You observed 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 1000MB. How would you tune your `io.sort.mb` value to achieve maximum memory to disk I/O ratio?

- A. For a 1GB child heap size an `io.sort.mb` of 128 MB will always maximize memory to disk I/O
- B. Increase the `io.sort.mb` to 1GB
- C. Decrease the `io.sort.mb` value to 0
- 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 3

Assume you have a file named `foo.txt` in your local directory. You issue the following three commands:

```
Hadoop fs mkdir input
```

```
Hadoop fs put foo.txt input/foo.txt
```

```
Hadoop fs put foo.txt input
```

What happens when you issue the third command?

- A. The write succeeds, overwriting `foo.txt` in HDFS with no warning
- B. The file is uploaded and stored as a plain file named `input`
- C. You get a warning that `foo.txt` is being overwritten

- D. You get an error message telling you that foo.txt already exists, and asking you if you would like to overwrite it.
- E. You get a error message telling you that foo.txt already exists. The file is not written to HDFS
- F. You get an error message telling you that input is not a directory
- G. The write silently fails

Correct Answer: CE

---

#### QUESTION 4

You are migrating a cluster from MApReduce version 1 (MRv1) to MapReduce version 2 (MRv2) on YARN. You want to maintain your MRv1 TaskTracker slot capacities when you migrate. What should you do/

- A. Configure yarn.applicationmaster.resource.memory-mb and yarn.applicationmaster.resource.cpu-vcores so that ApplicationMaster container allocations match the capacity you require.
- B. You don't need to configure or balance these properties in YARN as YARN dynamically balances resource management capabilities on your cluster
- C. Configure mapred.tasktracker.map.tasks.maximum and mapred.tasktracker.reduce.tasks.maximum ub yarn-site.xml to match your cluster's capacity set by the yarn-scheduler.minimum-allocation
- D. Configure yarn.nodemanager.resource.memory-mb and yarn.nodemanager.resource.cpu- vcores to match the capacity you require under YARN for each NodeManager

Correct Answer: D

---

#### QUESTION 5

Your Hadoop cluster contains nodes in three racks. You have not configured the dfs.hosts property in the NameNode's configuration file. What results?

- A. The NameNode will update the dfs.hosts property to include machines running the DataNode daemon on the next NameNode reboot or with the command dfsadmin refreshNodes
- B. No new nodes can be added to the cluster until you specify them in the dfs.hosts file
- C. Any machine running the DataNode daemon can immediately join the cluster
- D. Presented with a blank dfs.hosts property, the NameNode will permit DataNodes specified in mapred.hosts to join the cluster

Correct Answer: C