



# 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.geekcert.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

In CDH4 and later, which file contains a serialized form of all the directory and files inodes in the filesystem, giving the NameNode a persistent checkpoint of the filesystem metadata?

- A. fstime
- B. VERSION
- C. Fsimage\_N (where N reflects transactions up to transaction ID N)
- D. Edits\_N-M (where N-M transactions between transaction ID N and transaction ID N)

Correct Answer: C

---

### QUESTION 2

Each node in your Hadoop cluster, running YARN, has 64GB memory and 24 cores. Your yarn.site.xml has the following configuration:

```
yarn.nodemanager.resource.memory-mb  
32768
```

```
yarn.nodemanager.resource.cpu-vcores  
12
```

You want YARN to launch no more than 16 containers per node. What should you do?

- A. Modify yarn-site.xml with the following property: yarn.scheduler.minimum-allocation-mb 2048
- B. Modify yarn-sites.xml with the following property: yarn.scheduler.minimum-allocation-mb 4096
- C. Modify yarn-site.xml with the following property: yarn.nodemanager.resource.cpu-vcores
- D. No action is needed: YARN's dynamic resource allocation automatically optimizes the node memory and cores

Correct Answer: A

---

### QUESTION 3



You decide to create a cluster which runs HDFS in High Availability mode with automatic failover, using Quorum Storage. What is the purpose of ZooKeeper in such a configuration?

- A. It only keeps track of which NameNode is Active at any given time
- B. It monitors an NFS mount point and reports if the mount point disappears
- C. It both keeps track of which NameNode is Active at any given time, and manages the Edits file. Which is a log of changes to the HDFS filesystem
- D. It only manages the Edits file, which is log of changes to the HDFS filesystem
- E. Clients connect to ZooKeeper to determine which NameNode is Active

Correct Answer: A

#### QUESTION 4

You are running a Hadoop cluster with MapReduce version 2 (MRv2) on YARN. You consistently see that MapReduce map tasks on your cluster are running slowly because of excessive garbage collection of JVM, how do you increase JVM heap size property to 3GB to optimize performance?

- A. `yarn.application.child.java.opts=-Xsx3072m`
- B. `yarn.application.child.java.opts=-Xmx3072m`
- C. `mapreduce.map.java.opts=-Xms3072m`
- D. `mapreduce.map.java.opts=-Xmx3072m`

Correct Answer: C

#### QUESTION 5

A user comes to you, complaining that when she attempts to submit a Hadoop job, it fails. There is a Directory in HDFS named `/data/input`. The Jar is named `j.jar`, and the driver class is named `DriverClass`.

She runs the command:

```
Hadoop jar j.jar DriverClass /data/input/data/output
```

The error message returned includes the line: `PrivilegedActionException as:training (auth:SIMPLE) cause:org.apache.hadoop.mapreduce.lib.input.InvalidInputException: Input path does not exist: file:/data/input`

What is the cause of the error?

- A. The user is not authorized to run the job on the cluster
- B. The output directory already exists
- C. The name of the driver has been spelled incorrectly on the command line
- D. The directory name is misspelled in HDFS



E. The Hadoop configuration files on the client do not point to the cluster

Correct Answer: A

[CCA-500 PDF Dumps](#)

[CCA-500 VCE Dumps](#)

[CCA-500 Braindumps](#)