

# 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.leads4pass.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**

Cluster Summary:

45 files and directories, 12 blocks = 57 total. Heap size is 15.31 MB/193.38MB(7%)

Configured capacity	:	17.33GB
DFS Used	:	144KB
Non DFS Used	:	5.49GB
DFS Remaining	:	11.84GB
DFS Used %	:	0%
DFS Remaining %	:	68.32GB
Live Nodes	:	6
Dead Nodes	:	1
Decommissioning Nodes	:	0
Number of Under-Replicated Blocks	:	6

Refer to the above screenshot.

You configure a Hadoop cluster with seven DataNodes and on of your monitoring UIs displays the details shown in the exhibit.

What does the this tell you?

- A. The DataNode JVM on one host is not active
- B. Because your under-replicated blocks count matches the Live Nodes, one node is dead, and your DFS Used % equals 0%, you can't be certain that your cluster has all the data you've written it.
- C. Your cluster has lost all HDFS data which had bocks stored on the dead DatNode
- D. The HDFS cluster is in safe mode

Correct Answer: A

---

**QUESTION 2**

You are running a Hadoop cluster with a NameNode on host mynamenode, a secondary NameNode on host mysecondarynamenode and several DataNodes.

Which best describes how you determine when the last checkpoint happened?

- A. Execute hdfs namenode report on the command line and look at the Last Checkpoint information
- B. Execute hdfs dfsadmin saveNamespace on the command line which returns to you the last checkpoint value in fstime file

C. Connect to the web UI of the Secondary NameNode (<http://mysecondary:50090/>) and look at the "Last Checkpoint" information

D. Connect to the web UI of the NameNode (<http://mynamenode:50070/>) and look at the "Last Checkpoint" information

Correct Answer: C

---

### QUESTION 3

During the execution of a MapReduce v2 (MRv2) job on YARN, where does the Mapper place the intermediate data of each Map Task?

A. The Mapper stores the intermediate data on the node running the Job's ApplicationMaster so that it is available to YARN ShuffleService before the data is presented to the Reducer

B. The Mapper stores the intermediate data in HDFS on the node where the Map tasks ran in the HDFS / usercache/and(user)/apache/application\_and(appid) directory for the user who ran the job

C. The Mapper transfers the intermediate data immediately to the reducers as it is generated by the Map Task

D. YARN holds the intermediate data in the NodeManager's memory (a container) until it is transferred to the Reducer

E. The Mapper stores the intermediate data on the underlying filesystem of the local disk in the directories yarn.nodemanager.local-DIFS

Correct Answer: E

---

### QUESTION 4

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 5

You have a cluster running with the fair Scheduler enabled. There are currently no jobs running on the cluster, and you submit a job A, so that only job A is running on the cluster. A while later, you submit Job B. now Job A and Job B are running on the cluster at the same time. How will the Fair Scheduler handle these two jobs? (Choose two)

- A. When Job B gets submitted, it will get assigned tasks, while job A continues to run with fewer tasks.
- B. When Job B gets submitted, Job A has to finish first, before job B can gets scheduled.
- C. When Job A gets submitted, it doesn't consumes all the task slots.
- D. When Job A gets submitted, it consumes all the task slots.

Correct Answer: B

---

## QUESTION 6

Which process instantiates user code, and executes map and reduce tasks on a cluster running MapReduce v2 (MRv2) on YARN?

- A. NodeManager
- B. ApplicationMaster
- C. TaskTracker
- D. JobTracker
- E. NameNode
- F. DataNode
- G. ResourceManager

Correct Answer: A

---

## QUESTION 7

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 8

On a cluster running MapReduce v2 (MRv2) on YARN, a MapReduce job is given a directory of 10 plain text files as its input directory. Each file is made up of 3 HDFS blocks. How many Mappers will run?

- A. We cannot say; the number of Mappers is determined by the ResourceManager
- B. We cannot say; the number of Mappers is determined by the developer
- C. 30
- D. 3
- E. 10
- F. We cannot say; the number of mappers is determined by the ApplicationMaster

Correct Answer: E

---

## QUESTION 9

You have a cluster running with a FIFO scheduler enabled. You submit a large job A to the cluster, which you expect to run for one hour. Then, you submit job B to the cluster, which you expect to run a couple of minutes only.

You submit both jobs with the same priority.

Which two best describes how FIFO Scheduler arbitrates the cluster resources for job and its tasks?

(Choose two)

- A. Because there is a more than a single job on the cluster, the FIFO Scheduler will enforce a limit on the percentage of resources allocated to a particular job at any given time
- B. Tasks are scheduled on the order of their job submission
- C. The order of execution of job may vary
- D. Given job A and submitted in that order, all tasks from job A are guaranteed to finish before all tasks from job B
- E. The FIFO Scheduler will give, on average, and equal share of the cluster resources over the job lifecycle
- F. The FIFO Scheduler will pass an exception back to the client when Job B is submitted, since all slots on the cluster are use

Correct Answer: AD

---

## QUESTION 10

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-vccores
- D. No action is needed: YARN's dynamic resource allocation automatically optimizes the node memory and cores

Correct Answer: A

[Latest CCA-500 Dumps](#)

[CCA-500 Practice Test](#)

[CCA-500 Braindumps](#)