



CCA-505^{Q&As}

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.pass4itsure.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 are upgrading a Hadoop cluster from HDFS and MapReduce version 1 (MRv1) to one running HDFS and MapReduce version 2 (MRv2) on YARN. You want to set and enforce a block of 128MB for all new files written to the cluster after the upgrade. What should you do?

- A. Set `dfs.block.size` to 128M on all the worker nodes, on all client machines, and on the NameNode, and set the parameter to final.
- B. Set `dfs.block.size` to 134217728 on all the worker nodes, on all client machines, and on the NameNode, and set the parameter to final.
- C. Set `dfs.block.size` to 134217728 on all the worker nodes and client machines, and set the parameter to final. You do need to set this value on the NameNode.
- D. Set `dfs.block.size` to 128M on all the worker nodes and client machines, and set the parameter to final. You do need to set this value on the NameNode.
- E. You cannot enforce this, since client code can always override this value.

Correct Answer: C

QUESTION 2

You are configuring a cluster running HDFS, MapReduce version 2 (MRv2) on YARN running Linux. How must you format the underlying filesystem of each DataNode?

- A. They must not be formatted - - HDFS will format the filesystem automatically
- B. They may be formatted in any Linux filesystem
- C. They must be formatted as HDFS
- D. They must be formatted as either ext3 or ext4

Correct Answer: D

QUESTION 3

You are configuring your cluster to run HDFS and MapReduce v2 (MRv2) on YARN. Which daemons need to be installed on your clusters master nodes? (Choose Two)

- A. ResourceManager
- B. DataNode
- C. NameNode
- D. JobTracker
- E. TaskTracker



F. HMaster

Correct Answer: AC

QUESTION 4

On a cluster running CDH 5.0 or above, you use the `hadoop fs put` command to write a 300MB file into a previously empty directory using an HDFS block of 64MB. Just after this command has finished writing 200MB of this file, what would another user see when they look in the directory?

- A. They will see the file with its original name. If they attempt to view the file, they will get a `ConcurrentFileAccessException` until the entire file write is completed on the cluster
- B. They will see the file with a `._COPYING_` extension on its name. If they attempt to view the file, they will get a `ConcurrentFileAccessException` until the entire file write is completed on the cluster.
- C. They will see the file with a `._COPYING_` extension on its name. If they view the file, they will see contents of the file up to the last completed block (as each 64MB block is written, that block becomes available)
- D. The directory will appear to be empty until the entire file write is completed on the cluster

Correct Answer: C

QUESTION 5

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 at least 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

[CCA-505 PDF Dumps](#)

[CCA-505 Study Guide](#)

[CCA-505 Exam Questions](#)