



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.geekcert.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 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 3

Assuming a cluster running HDFS, MapReduce version 2 (MRv2) on YARN with all settings at their default, what do you need to do when adding a new slave node to a cluster?

- A. Nothing, other than ensuring that DNS (or `/etc/hosts` files on all machines) contains an entry for the new node.
- B. Restart the NameNode and ResourceManager daemons and resubmit any running jobs
- C. Increase the value of `dfs.number.of.replicas` in `hdfs-site.xml`
- D. Add a new entry to `/etc/nodes` on the NameNode host.



E. Restart the NameNode daemon.

Correct Answer: B

QUESTION 4

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. ResourceManager
- D. TaskTracker
- E. JobTracker
- F. DataNode
- G. NameNode

Correct Answer: E

QUESTION 5

What processes must you do if you are running a Hadoop cluster with a single NameNode and six DataNodes, and you want to change a configuration parameter so that it affects all six DataNodes.

- A. You must modify the configuration file on each of the six DataNode machines.
- B. You must restart the NameNode daemon to apply the changes to the cluster
- C. You must restart all six DataNode daemon to apply the changes to the cluster
- D. You don't need to restart any daemon, as they will pick up changes automatically
- E. You must modify the configuration files on the NameNode only. DataNodes read their configuration from the master nodes.

Correct Answer: BE

[CCA-505 VCE Dumps](#)

[CCA-505 Practice Test](#)

[CCA-505 Exam Questions](#)