



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

For each YARN Job, the Hadoop framework generates task log files. Where are Hadoop's files stored?

- A. In HDFS, In the directory of the user who generates the job
- B. On the local disk of the slave node running the task
- C. Cached In the YARN container running the task, then copied into HDFS on job completion
- D. Cached by the NodeManager managing the job containers, then written to a log directory on the NameNode

Correct Answer: B

QUESTION 3

Your cluster has the following characteristics:

A rack aware topology is configured and on

Replication is not set to 3

Cluster block size is set to 64 MB

Which describes the file read process when a client application connects into the cluster and requests a 50MB file?

- A. The client queries the NameNode which retrieves the block from the nearest DataNode to the client and then passes that block back to the client.



B. The client queries the NameNode for the locations of the block, and reads from a random location in the list it retrieves to eliminate network I/O leads by balancing which nodes it retrieves data from at any given time.

C. The client queries the NameNode for the locations of the block, and reads all three copies. The first copy to complete transfer to the client is the one the client reads as part of Hadoop's

speculative execution framework.

D. The client queries the NameNode for the locations of the block, and reads from the first location in the list it receives.

Correct Answer: A

QUESTION 4

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 5

Which YARN daemon or service negotiates map and reduce Containers from the Scheduler, tracking their status and monitoring for progress?

A. ResourceManager

B. ApplicationMaster

C. NodeManager

D. ApplicationManager

Correct Answer: B

[Latest CCA-505 Dumps](#)

[CCA-505 PDF Dumps](#)

[CCA-505 Exam Questions](#)