



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.passapply.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

On a cluster running MapReduce v2 (MRv2) on YARN, a MapReduce job is given a directory of 10 plain text 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 ApplicationManager
- C. We cannot say; the number of Mappers is determined by the developer
- D. 30
- E. 3
- F. 10

Correct Answer: E

QUESTION 2

You have a 20 node Hadoop cluster, with 18 slave nodes and 2 master nodes running HDFS High Availability (HA). You want to minimize the chance of data loss in your cluster. What should you do?

- A. Add another master node to increase the number of nodes running the JournalNode which increases the number of machines available to HA to create a quorum
- B. Configure the cluster's disk drives with an appropriate fault tolerant RAID level
- C. Run the ResourceManager on a different master from the NameNode in the order to load share HDFS metadata processing
- D. Run a Secondary NameNode on a different master from the NameNode in order to load provide automatic recovery from a NameNode failure
- E. Set an HDFS replication factor that provides data redundancy, protecting against failure

Correct Answer: C

QUESTION 3

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 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

You have installed a cluster running HDFS and MapReduce version 2 (MRv2) on YARN. You have no afs.hosts entry()ies in your hdfs-alte.xml configuration file. You configure a new worker node by setting fs.default.name in its configuration files to point to the NameNode on your cluster, and you start the DataNode daemon on that worker node.

What do you have to do on the cluster to allow the worker node to join, and start storing HDFS blocks?

- A. Nothing; the worker node will automatically join the cluster when the DataNode daemon is started.
- B. Without creating a dfs.hosts file or making any entries, run the command `hadoop dfsadmin refreshHadoop` on the NameNode
- C. Create a dfs.hosts file on the NameNode, add the worker node\\'s name to it, then issue the command `hadoop dfsadmin refreshNodes` on the NameNode
- D. Restart the NameNode

Correct Answer: B

[Latest CCA-505 Dumps](#)

[CCA-505 VCE Dumps](#)

[CCA-505 Braindumps](#)