# 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 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

**QUESTION 2**

Which three basic configuration parameters must you set to migrate your cluster from MapReduce1 (MRv1) to MapReduce v2 (MRv2)?

A. Configure the NodeManager hostname and enable services on YARN by setting the following property in yarn-site.xml: yarn.nodemanager.hostname your_nodeManager_hostname

B. Configure the number of map tasks per job on YARN by setting the following property in mapredsite.xml: mapreduce.job.maps 2

C. Configure MapReduce as a framework running on YARN by setting the following property in mapredsite.xml: mapreduce.framework.name yarn

D. Configure the ResourceManager hostname and enable node services on YARN by setting the following property in yarn-site.xml: yarn.resourcemanager.hostname your_responseManager_hostname

E. Configure a default scheduler to run on YARN by setting the following property in sapred- site.xml: mapreduce.jobtracker.taskScheduler org.apache.hadoop.mapred.JobQueueTaskScheduler

F. Configure the NodeManager to enable MapReduce services on YARN by adding following property in yarn-site.xml: yarn.nodemanager.aux-services mapreduce_shuffle

Correct Answer: ABD

**QUESTION 3**

Your Hadoop cluster contains nodes in three racks. You have NOT configured the dfs.hosts property in the NameNode\\'s configuration file. What results?

A. No new nodes can be added to the cluster until you specify them in the dfs.hosts file

B. Presented with a blank dfs.hosts property, the NameNode will permit DatNode specified in mapred.hosts to join the cluster

C. Any machine running the DataNode daemon can immediately join the cluster

D. The NameNode will update the dfs.hosts property to include machine running DataNode daemon on the next NameNode reboot or with the command dfsadmin -refreshNodes

Correct Answer: C

**QUESTION 4**

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 you 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 5**

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

A. The Mapper stores the intermediate data on the mode running the job\\'s ApplicationMaster so that is available to YARN\\'s 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]sppcache/application_and(appid) directory for the user who ran the job

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

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

E. The Mapper transfers the intermediate data immediately to the Reducers as it generated by the Map task

Correct Answer: D