

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.pass2lead.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

Your cluster implements HDFS High Availability (HA). Your two NameNodes are named nn01 and nn02. What occurs when you execute the command: `hdfs haadmin failover nn01 nn02`

- A. nn02 becomes the standby NameNode and nn01 becomes the active NameNode
- B. nn02 is fenced, and nn01 becomes the active NameNode
- C. nn01 becomes the standby NameNode and nn02 becomes the active NameNode
- D. nn01 is fenced, and nn02 becomes the active NameNode

Correct Answer: D

QUESTION 2

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 daemons 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

QUESTION 3

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

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

Correct Answer: AC

QUESTION 4

You have installed a cluster running HDFS and MapReduce version 2 (MRv2) on YARN. You have no `dfs.hosts` entry(ies) in your `hdfs-site.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 5

You are planning a Hadoop cluster and considering implementing 10 Gigabit Ethernet as the network fabric. Which workloads benefit the most from a faster network fabric?

- A. When your workload generates a large amount of output data, significantly larger than amount of intermediate data
- B. When your workload generates a large amount of intermediate data, on the order of the input data itself
- C. When workload consumers a large amount of input data, relative to the entire capacity of HDFS
- D. When your workload consists of processor-intensive tasks

Correct Answer: B

[CCA-505 PDF Dumps](#)

[CCA-505 VCE Dumps](#)

[CCA-505 Braindumps](#)