



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 a Hadoop cluster running HDFS, and a gateway machine external to the cluster from which clients submit jobs. What do you need to do in order to run on the cluster and submit jobs from the command line of the gateway machine?

- A. Install the implad daemon, statestored daemon, and catalogd daemon on each machine in the cluster and on the gateway node
- B. Install the impalad daemon on each machine in the cluster, the statestored daemon and catalogd daemon on one machine in the cluster, and the impala shell on your gateway machine
- C. Install the impalad daemon and the impala shell on your gateway machine, and the statestored daemon and catalog daemon on one of the nodes in the cluster
- D. Install the impalad daemon, the statestored daemon, the catalogd daemon, and the impala shell on your gateway machine
- E. Install the impalad daemon, statestored daemon, and catalogd daemon on each machine in the cluster, and the impala shell on your gateway machine

Correct Answer: B

QUESTION 2

You are the hadoop fs put command to add a file "sales.txt" to HDFS. This file is small enough that it fits into a single block, which is replicated to three nodes in your cluster (with a replication factor of 3). One of the nodes holding this file (a single block) fails. How will the cluster handle the replication of this file in this situation/

- A. The cluster will re-replicate the file the next time the system administrator reboots the NameNode daemon (as long as the file's replication doesn't fall two)
- B. This file will be immediately re-replicated and all other HDFS operations on the cluster will halt until the cluster's replication values are restored
- C. The file will remain under-replicated until the administrator brings that nodes back online
- D. The file will be re-replicated automatically after the NameNode determines it is under replicated based on the block reports it receives from the DataNodes

Correct Answer: B

QUESTION 3

A slave node in your cluster has four 2TB hard drives installed (4 x 2TB). The DataNode is configured to store HDFS blocks on the disks. You set the value of the dfs.datanode.du.reserved parameter to 100GB. How does this alter HDFS block storage?

- A. A maximum of 100 GB on each hard drive may be used to store HDFS blocks
- B. All hard drives may be used to store HDFS blocks as long as atleast 100 GB in total is available on the node



- C. 100 GB on each hard drive may not be used to store HDFS blocks
- D. 25 GB on each hard drive may not be used to store HDFS blocks

Correct Answer: B

QUESTION 4

Your Hadoop cluster is configured with HDFS and MapReduce version 2 (MRv2) on YARN. Can you configure a worker node to run a NodeManager daemon but not a DataNode daemon and still have a function cluster?

- A. Yes. The daemon will receive data from the NameNode to run Map tasks
- B. Yes. The daemon will get data from another (non-local) DataNode to run Map tasks
- C. Yes. The daemon will receive Reduce tasks only

Correct Answer: A

QUESTION 5

You are working on a project where you need to chain together MapReduce, Pig jobs. You also needs the ability to use forks, decision, and path joins. Which ecosystem project should you use to perform these actions?

- A. Oozie
- B. Zookeeper
- C. HBase
- D. Sqoop
- E. HUE

Correct Answer: A

[CCA-505 PDF Dumps](#)

[CCA-505 Practice Test](#)

[CCA-505 Study Guide](#)