



# CCA-505<sup>Q&As</sup>

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 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 2

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 3

You are configuring a cluster running HDFS, MapReduce version 2 (MRv2) on YARN running Linux. How must you format the underlying filesystem of each DataNode?

- A. They must not formatted - - HDFS will format the filesystem automatically
- B. They may be formatted in any Linux filesystem
- C. They must be formatted as HDFS
- D. They must be formatted as either ext3 or ext4

Correct Answer: D

---



#### QUESTION 4

You decide to create a cluster which runs HDFS in High Availability mode with automatic failover, using Quorum-based Storage. What is the purpose of ZooKeeper in such a configuration?

- A. It manages the Edits file, which is a log changes to the HDFS filesystem.
- B. It monitors an NFS mount point and reports if the mount point disappears
- C. It both keeps track of which NameNode is Active at any given time, and manages the Edits file, which is a log of changes to the HDFS filesystem
- D. It only keeps track of which NameNode is Active at any given time
- E. Clients connect to ZooKeeper to determine which NameNode is Active

Correct Answer: D

---

#### QUESTION 5

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

[CCA-505 PDF Dumps](#)

[CCA-505 Study Guide](#)

[CCA-505 Exam Questions](#)