

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

Cloudera Certified Administrator for Apache Hadoop (CCA) CDH5  
Upgrade Exam

**Pass Cloudera CCA-505 Exam with 100% Guarantee**

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.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

Which is the default scheduler in YARN?

- A. Fair Scheduler
- B. FIFO Scheduler
- C. Capacity Scheduler
- D. YARN doesn't configure a default scheduler. You must first assign an appropriate scheduler class in yarn-site.xml

Correct Answer: C

---

## QUESTION 2

You are working on a project where you need to chain together MapReduce, Pig jobs. You also need 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

---

## QUESTION 3

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 below 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 node 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 4

On a cluster running CDH 5.0 or above, you use the `hadoop fs put` command to write a 300MB file into a previously empty directory using an HDFS block of 64MB. Just after this command has finished writing 200MB of this file, what would another user see when they look in the directory?

- A. They will see the file with its original name. If they attempt to view the file, they will get a `ConcurrentFileAccessException` until the entire file write is completed on the cluster
- B. They will see the file with a `._COPYING_` extension on its name. If they attempt to view the file, they will get a `ConcurrentFileAccessException` until the entire file write is completed on the cluster.
- C. They will see the file with a `._COPYING_` extension on its name. If they view the file, they will see contents of the file up to the last completed block (as each 64MB block is written, that block becomes available)
- D. The directory will appear to be empty until the entire file write is completed on the cluster

Correct Answer: C

---

## QUESTION 5

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

[CCA-505 PDF Dumps](#)

[CCA-505 Practice Test](#)

[CCA-505 Study Guide](#)