



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

Which YARN process runs as "controller O" of a submitted job and is responsible for resource requests?

- A. ResourceManager
- B. NodeManager
- C. JobHistoryServer
- D. ApplicationMaster
- E. JobTracker
- F. ApplicationManager

Correct Answer: D

QUESTION 2

Assuming a cluster running HDFS, MapReduce version 2 (MRv2) on YARN with all settings at their default, what do you need to do when adding a new slave node to a cluster?

- A. Nothing, other than ensuring that DNS (or /etc/hosts files on all machines) contains an entry for the new node.
- B. Restart the NameNode and ResourceManager daemons and resubmit any running jobs
- C. Increase the value of dfs.number.of.replicas in hdfs-site.xml
- D. Add a new entry to /etc/nodes on the NameNode host.
- E. Restart the NameNode daemon.

Correct Answer: B

QUESTION 3

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 4

You observe that the number of spilled records from Map tasks far exceeds the number of map output records. Your child heap size is 1GB and your io.sort.mb value is set to 100 MB. How would you tune your io.sort.mb value to achieve maximum memory to disk I/O ratio?

- A. Decrease the io.sort.mb value to 0
- B. Increase the io.sort.mb to 1GB
- C. For 1GB child heap size an io.sort.mb of 128 MB will always maximize memory to disk I/O
- D. Tune the io.sort.mb value until you observe that the number of spilled records equals (or is as close to equals) the number of map output records

Correct Answer: D

QUESTION 5

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

[CCA-505 Practice Test](#)

[CCA-505 Exam Questions](#)

[CCA-505 Braindumps](#)