

CCA-505^{Q&As}

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

You want a node to only swap Hadoop daemon data from RAM to disk when absolutely necessary. What should you do?

- A. Delete the /swapfile file on the node
- B. Set vm.swappiness to 0 in /etc/sysctl.conf
- C. Set the ram.swap parameter to 0 in core-site.xml
- D. Delete the /etc/swap file on the node
- E. Delete the /dev/vmswap file on the node

Correct Answer: B

QUESTION 2

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 3

In CDH4 and later, which file contains a serialized form of all the directory and files inodes in the filesystem, giving the NameNode a persistent checkpoint of the filesystem metadata?

- A. fstime
- B. VERSION
- C. Fsimage_N (Where N reflects all transactions up to transaction ID N)
- D. Edits_N-M (Where N-M specifies transactions between transactions ID N and transaction ID N)

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 PDF Dumps](#)

[CCA-505 VCE Dumps](#)

[CCA-505 Practice Test](#)