

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.pass2lead.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

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 2

On a cluster running MapReduce v2 (MRv2) on YARN, a MapReduce job is given a directory of 10 plain text as its input directory. Each file is made up of 3 HDFS blocks. How many Mappers will run?

- A. We cannot say; the number of Mappers is determined by the ResourceManager
- B. We cannot say; the number of Mappers is determined by the ApplicationManager
- C. We cannot say; the number of Mappers is determined by the developer
- D. 30
- E. 3
- F. 10

Correct Answer: E

QUESTION 3

You want to understand more about how users browse your public website. For example, you want to know which pages they visit prior to placing an order. You have a server farm of 200 web servers hosting your website. Which is the most efficient process to gather these web server logs into your Hadoop cluster for analysis?

- A. Sample the web server logs from web servers and copy them into HDFS using curl
- B. Ingest the server web logs into HDFS using Flume
- C. Import all users clicks from your OLTP databases into Hadoop using Sqoop
- D. Write a MapReduce job with the web servers as mappers and the Hadoop cluster nodes as reducers
- E. Channel these clickstreams into Hadoop using Hadoop Streaming

Correct Answer: AB

QUESTION 4

You are planning a Hadoop cluster and considering implementing 10 Gigabit Ethernet as the network fabric. Which workloads benefit the most from a faster network fabric?

- A. When your workload generates a large amount of output data, significantly larger than amount of intermediate data
- B. When your workload generates a large amount of intermediate data, on the order of the input data itself
- C. When workload consumers a large amount of input data, relative to the entire capacity of HDFS
- D. When your workload consists of processor-intensive tasks

Correct Answer: B

QUESTION 5

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

[Latest CCA-505 Dumps](#)

[CCA-505 PDF Dumps](#)

[CCA-505 Exam Questions](#)