



CCA-500^{Q&As}

Cloudera Certified Administrator for Apache Hadoop (CCA-H)

Pass Cloudera CCA-500 Exam with 100% Guarantee

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

<https://www.pass4itsure.com/cca-500.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 decide to create a cluster which runs HDFS in High Availability mode with automatic failover, using Quorum Storage. What is the purpose of ZooKeeper in such a configuration?

- A. It only keeps track of which NameNode is Active at any given time
- 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 manages the Edits file, which is log of changes to the HDFS filesystem
- E. Clients connect to ZooKeeper to determine which NameNode is Active

Correct Answer: A

QUESTION 2

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 transactions up to transaction ID N)
- D. Edits_N-M (where N-M transactions between transaction ID N and transaction ID N)

Correct Answer: C

QUESTION 3

What two processes must you do if you are running a Hadoop cluster with a single NameNode and six DataNodes, and you want to change a configuration parameter so that it affects all six DataNodes. (Choose two)

- A. You must modify the configuration files on the NameNode only. DataNodes read their configuration from the master nodes
- B. You must modify the configuration files on each of the DataNodes machines
- C. You don't need to restart any daemon, as they will pick up changes automatically
- D. You must restart the NameNode daemon to apply the changes to the cluster
- E. You must restart all six DataNode daemon to apply the changes to the cluster

Correct Answer: AD

**QUESTION 4**

You are migrating a cluster from MapReduce version 1 (MRv1) to MapReduce version 2 (MRv2) on YARN. You want to maintain your MRv1 TaskTracker slot capacities when you migrate. What should you do/

- A. Configure `yarn.applicationmaster.resource.memory-mb` and `yarn.applicationmaster.resource.cpu-vcores` so that ApplicationMaster container allocations match the capacity you require.
- B. You don't need to configure or balance these properties in YARN as YARN dynamically balances resource management capabilities on your cluster
- C. Configure `mapred.tasktracker.map.tasks.maximum` and `mapred.tasktracker.reduce.tasks.maximum` in `yarn-site.xml` to match your cluster's capacity set by the `yarn-scheduler.minimum-allocation`
- D. Configure `yarn.nodemanager.resource.memory-mb` and `yarn.nodemanager.resource.cpu-vcores` to match the capacity you require under YARN for each NodeManager

Correct Answer: D

QUESTION 5

Which three basic configuration parameters must you set to migrate your cluster from MapReduce 1 (MRv1) to MapReduce V2 (MRv2)? (Choose three)

- A. Configure the NodeManager to enable MapReduce services on YARN by setting the following property in `yarn-site.xml`: `yarn.nodemanager.hostname your_nodeManager_shuffle`
- B. Configure the NodeManager hostname and enable node services on YARN by setting the following property in `yarn-site.xml`: `yarn.nodemanager.hostname your_nodeManager_hostname`
- C. Configure a default scheduler to run on YARN by setting the following property in `mapred-site.xml`: `mapreduce.jobtracker.taskScheduler org.apache.hadoop.mapred.JobQueueTaskScheduler`
- D. Configure the number of map tasks per job on YARN by setting the following property in `mapred-site.xml`: `mapreduce.job.maps 2`
- E. Configure the ResourceManager hostname and enable node services on YARN by setting the following property in `yarn-site.xml`: `yarn.resourcemanager.hostname your_resourceManager_hostname`
- F. Configure MapReduce as a Framework running on YARN by setting the following property in `mapred-site.xml`: `mapreduce.framework.name yarn`

Correct Answer: AEF

[Latest CCA-500 Dumps](#)

[CCA-500 PDF Dumps](#)

[CCA-500 VCE Dumps](#)