



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

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 2

Identify two features/issues that YARN is designed to address: (Choose two)

- A. Standardize on a single MapReduce API
- B. Single point of failure in the NameNode
- C. Reduce complexity of the MapReduce APIs
- D. Resource pressure on the JobTracker
- E. Ability to run framework other than MapReduce, such as MPI
- F. HDFS latency

Correct Answer: DE

QUESTION 3

Which YARN process runs as "container 0" of a submitted job and is responsible for resource requests?

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



Correct Answer: C

QUESTION 4

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. Set an HDFS replication factor that provides data redundancy, protecting against node failure
- C. Run a Secondary NameNode on a different master from the NameNode in order to provide automatic recovery from a NameNode failure.
- D. Run the ResourceManager on a different master from the NameNode in order to load-share HDFS metadata processing
- E. Configure the cluster's disk drives with an appropriate fault tolerant RAID level

Correct Answer: D

QUESTION 5

You are running Hadoop cluster with all monitoring facilities properly configured.

Which scenario will go undetected?

- A. HDFS is almost full
- B. The NameNode goes down
- C. A DataNode is disconnected from the cluster
- D. Map or reduce tasks that are stuck in an infinite loop
- E. MapReduce jobs are causing excessive memory swaps

Correct Answer: B

[Latest CCA-500 Dumps](#)

[CCA-500 PDF Dumps](#)

[CCA-500 Exam Questions](#)