

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.leads4pass.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

Given:

```
[user1@host1 ~] yarn application -llst
```

Total Applications: 3

Application ID	Application-Name	Application-Type	User	Queue	State	Final-State	Progress
Application_1374638600275_0109	Sleep Job	MAPREDUCE	user1	KILLED	KILLED	KILLED	100%
Application_1374638600275_0121	Sleep Job	MAPREDUCE	user1	FINISHED	SUCCEEDED	SUCCEEDED	100%
Application_1374638600275_0020	Sleep Job	MAPREDUCE	user1	FINISHED	SUCCEEDED	SUCCEEDED	100%

You want to clean up this list by removing jobs where the State is KILLED. What command you enter?

- A. Yarn application refreshJobHistory
- B. Yarn application kill application_1374638600275_0109
- C. Yarn radmin refreshQueue
- D. Yarn radmin kill application_1374638600275_0109

Correct Answer: B

QUESTION 2

You have a Hadoop cluster HDFS, and a gateway machine external to the cluster from which clients submit jobs. What do you need to do in order to run Impala on the cluster and submit jobs from the command line of the gateway machine?

- A. Install the impalad daemon, statestored daemon, and daemon on each machine in the cluster, and the impala shell on your gateway machine
- B. Install the impalad daemon, the statestored daemon, the catalogd daemon, and the impala shell on your gateway machine
- C. Install the impalad daemon and the impala shell on your gateway machine, and the statestored daemon and catalogd daemon on one of the nodes in the cluster
- D. Install the impalad daemon on each machine in the cluster, the statestored daemon and catalogd daemon on one machine in the cluster, and the impala shell on your gateway machine
- E. Install the impalad daemon, statestored daemon, and catalogd daemon on each machine in the cluster and on the gateway node

Correct Answer: D

QUESTION 3

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

Which scenario will go undeselected?

- 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

QUESTION 4

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

QUESTION 5

Your cluster has the following characteristics:

A rack aware topology is configured and on

Replication is set to 3

Cluster block size is set to 64MB

Which describes the file read process when a client application connects into the cluster and requests a 50MB file?

- A. The client queries the NameNode for the locations of the block, and reads all three copies. The first copy to complete transfer to the client is the one the client reads as part of Hadoop's speculative execution framework.

- B. The client queries the NameNode for the locations of the block, and reads from the first location in the list it receives.
- C. The client queries the NameNode for the locations of the block, and reads from a random location in the list it receives to eliminate network I/O loads by balancing which nodes it retrieves data from any given time.
- D. The client queries the NameNode which retrieves the block from the nearest DataNode to the client then passes that block back to the client.

Correct Answer: B

[Latest CCA-500 Dumps](#)

[CCA-500 Practice Test](#)

[CCA-500 Study Guide](#)