



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.passapply.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

Your cluster's `mapred-start.xml` includes the following parameters

```
mapreduce.map.memory.mb 4096 mapreduce.reduce.memory.mb 8192
```

And any cluster's `yarn-site.xml` includes the following parameters

```
yarn.nodemanager.vmen-pmen-ratio 2.1
```

What is the maximum amount of virtual memory allocated for each map task before YARN will kill its Container?

- A. 4 GB
- B. 17.2 GB
- C. 8.9 GB
- D. 8.2 GB
- E. 24.6 GB

Correct Answer: D

QUESTION 2

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 3



Which command does Hadoop offer to discover missing or corrupt HDFS data?

- A. Hdfs fs du
- B. Hdfs fsck
- C. Dskchk
- D. The map-only checksum
- E. Hadoop does not provide any tools to discover missing or corrupt data; there is not need because three replicas are kept for each data block

Correct Answer: B

QUESTION 4

A user comes to you, complaining that when she attempts to submit a Hadoop job, it fails. There is a Directory in HDFS named /data/input. The Jar is named j.jar, and the driver class is named DriverClass.

She runs the command:

```
Hadoop jar j.jar DriverClass /data/input/data/output
```

The error message returned includes the line: PrivilegedActionException as:training (auth:SIMPLE) cause:org.apache.hadoop.mapreduce.lib.input.InvalidInputException: Input path does not exist: file:/data/input

What is the cause of the error?

- A. The user is not authorized to run the job on the cluster
- B. The output directory already exists
- C. The name of the driver has been spelled incorrectly on the command line
- D. The directory name is misspelled in HDFS
- E. The Hadoop configuration files on the client do not point to the cluster

Correct Answer: A

QUESTION 5

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

[CCA-500 PDF Dumps](#)

[CCA-500 Exam Questions](#)

[CCA-500 Braindumps](#)