



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**

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: 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 mapredsite.xml: mapreduce.framework.name yarn

Correct Answer: AEF

QUESTION 2

Your Hadoop cluster is configuring with HDFS and MapReduce version 2 (MRv2) on YARN. Can you configure a worker node to run a NodeManager daemon but not a DataNode daemon and still have a functional cluster?

- A. Yes. The daemon will receive data from the NameNode to run Map tasks
- B. Yes. The daemon will get data from another (non-local) DataNode to run Map tasks
- C. Yes. The daemon will receive Map tasks only
- D. Yes. The daemon will receive Reducer tasks only

Correct Answer: B

QUESTION 3

Which YARN daemon or service monitors a Controller's per-application resource using (e.g., memory CPU)?

- A. ApplicationMaster
- B. NodeManager
- C. ApplicationManagerService
- D. ResourceManager



Correct Answer: A

QUESTION 4

You have recently converted your Hadoop cluster from a MapReduce 1 (MRv1) architecture to MapReduce 2 (MRv2) on YARN architecture. Your developers are accustomed to specifying map and reduce tasks (resource allocation) tasks when they run jobs: A developer wants to know how specify to reduce tasks when a specific job runs. Which method should you tell that developers to implement?

- A. MapReduce version 2 (MRv2) on YARN abstracts resource allocation away from the idea of "tasks" into memory and virtual cores, thus eliminating the need for a developer to specify the number of reduce tasks, and indeed preventing the developer from specifying the number of reduce tasks.
- B. In YARN, resource allocations is a function of megabytes of memory in multiples of 1024mb. Thus, they should specify the amount of memory resource they need by executing `D mapreducereduces.memory-mb-2048`
- C. In YARN, the ApplicationMaster is responsible for requesting the resource required for a specific launch. Thus, executing `D yarn.applicationmaster.reduce.tasks=2` will specify that the ApplicationMaster launch two task contains on the worker nodes.
- D. Developers specify reduce tasks in the exact same way for both MapReduce version 1 (MRv1) and MapReduce version 2 (MRv2) on YARN. Thus, executing `D mapreduce.job.reduces-2` will specify reduce tasks.
- E. In YARN, resource allocation is function of virtual cores specified by the ApplicationManager making requests to the NodeManager where a reduce task is handled by a single container (and thus a single virtual core). Thus, the developer needs to specify the number of virtual cores to the NodeManager by executing `p yarn.nodemanager.cpu-vcores=2`

Correct Answer: D

QUESTION 5

What does CDH packaging do on install to facilitate Kerberos security setup?

- A. Automatically configures permissions for log files at `and MAPRED_LOG_DIR/userlogs`
- B. Creates users for hdfs and mapreduce to facilitate role assignment
- C. Creates directories for temp, hdfs, and mapreduce with the correct permissions
- D. Creates a set of pre-configured Kerberos keytab files and their permissions
- E. Creates and configures your kdc with default cluster values

Correct Answer: B

[Latest CCA-500 Dumps](#)

[CCA-500 Study Guide](#)

[CCA-500 Exam Questions](#)