



# CCA-505<sup>Q&As</sup>

Cloudera Certified Administrator for Apache Hadoop (CCAH) CDH5  
Upgrade Exam

## Pass Cloudera CCA-505 Exam with 100% Guarantee

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

<https://www.passapply.com/cca-505.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 is the default scheduler in YARN?

- A. Fair Scheduler
- B. FIFO Scheduler
- C. Capacity Scheduler
- D. YARN doesn't configure a default scheduler. You must first assign an appropriate scheduler class in yarn-site.xml

Correct Answer: C

---

### QUESTION 2

Your Hadoop cluster is configured 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 Reduce tasks only

Correct Answer: A

---

### QUESTION 3

You are migrating a cluster from MapReduce version 1 (MRv1) to MapReduce version 2 (MRv2) on YARN. To 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.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 yarn.nodemanager.resource.memory-mb and yarn.nodemanager.resource.cpu-vcores to match the capacity you require under YARN for each NodeManager
- D. Configure mapred.tasktracker.map.tasks.maximum and mapred.tasktracker.reduce.tasks.maximum in yarn.site.xml to match your cluster's configured capacity set by yarn.scheduler.minimum-allocation

Correct Answer: C

---

### QUESTION 4

You have converted your Hadoop cluster from a MapReduce 1 (MRv1) architecture to a 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 developer to implement?

- A. Developers specify reduce tasks in the exact same way for both MapReduce version 1 (MRv1) and MapReduce version 2 (MRv2) on YARN. Thus, executing `p mapreduce.job.reduce-2` will specify 2 reduce tasks.
- B. In YARN, the ApplicationMaster is responsible for requesting the resources required for a specific job. Thus, executing `p yarn.applicationmaster.reduce.tasks-2` will specify that the ApplicationMaster launch two task containers on the worker nodes.
- C. In YARN, resource allocation is a function of megabytes of memory in multiple of 1024mb. Thus, they should specify the amount of memory resource they need by executing `D mapreduce.reduce.memory-mp-2040`
- D. In YARN, resource allocation is a function of virtual cores specified by the ApplicationMaster making requests to the NodeManager where a reduce task is handled by a single container (and this 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`
- E. 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.

Correct Answer: D

---

#### QUESTION 5

Your cluster is configured with HDFS and MapReduce version 2 (MRv2) on YARN. What is the result when you execute: `hadoop jar samplejar.jar MyClass` on a client machine?

- A. SampleJar.jar is sent to the ApplicationMaster which allocation a container for Sample.jar
- B. SampleJar.Jar is serialized into an XML file which is submitted to the ApplicationMaster
- C. SampleJar.Jar is sent directly to the ResourceManager
- D. SampleJar.Jar is placed in a temporary directly in HDFS

Correct Answer: A

[Latest CCA-505 Dumps](#)

[CCA-505 PDF Dumps](#)

[CCA-505 Study Guide](#)