



## 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:

<http://www.CertBus.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
-  **80000+** Satisfied Customers



**Vendor:** Cloudera

**Exam Code:** CCA-500

**Exam Name:** Cloudera Certified Administrator for Apache Hadoop (CCA-H)

**Q&As:** Demo

### QUESTION 1

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

```
<name>mapreduce.map.memory.mb</name>
<value>4096</value>
<name>mapreduce.reduce.memory.mb</name>
<value>8192</value>
```

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

```
<name>yarn.nodemanager.vmen-pmen-ration</name>
<value>2.1</value>
```

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

Assuming you're not running HDFS Federation, what is the maximum number of NameNode daemons you should run on your cluster in order to avoid a "split-brain" scenario with your NameNode when running HDFS High Availability (HA) using Quorum-based storage?

- A. Two active NameNodes and two Standby NameNodes
- B. One active NameNode and one Standby NameNode
- C. Two active NameNodes and one Standby NameNode
- D. Unlimited. HDFS High Availability (HA) is designed to overcome limitations on the number of NameNodes you can deploy

**Correct Answer: B**

### QUESTION 3

Table schemas in Hive are:

- A. Stored as metadata on the NameNode
- B. Stored along with the data in HDFS
- C. Stored in the Metadata
- D. Stored in ZooKeeper

**Correct Answer: B**

### QUESTION 4

For each YARN job, the Hadoop framework generates task log file. Where are Hadoop task log files stored?

- A. Cached by the NodeManager managing the job containers, then written to a log directory on the NameNode
- B. Cached in the YARN container running the task, then copied into HDFS on job completion
- C. In HDFS, in the directory of the user who generates the job
- D. On the local disk of the slave node running the task

**Correct Answer:** D

**QUESTION 5**

You have a cluster running with the fair Scheduler enabled. There are currently no jobs running on the cluster, and you submit a job A, so that only job A is running on the cluster. A while later, you submit Job B. now Job A and Job B are running on the cluster at the same time. How will the Fair Scheduler handle these two jobs? (Choose two)

- A. When Job B gets submitted, it will get assigned tasks, while job A continues to run with fewer tasks.
- B. When Job B gets submitted, Job A has to finish first, before job B can get scheduled.
- C. When Job A gets submitted, it doesn't consume all the task slots.
- D. When Job A gets submitted, it consumes all the task slots.

**Correct Answer:** B

**QUESTION 6**

Each node in your Hadoop cluster, running YARN, has 64GB memory and 24 cores. Your yarn.site.xml has the following configuration:

```
<property>
<name>yarn.nodemanager.resource.memory-mb</name>
<value>32768</value>
</property>
<property>
<name>yarn.nodemanager.resource.cpu-vcores</name>
<value>12</value>
</property>
```

You want YARN to launch no more than 16 containers per node. What should you do?

- A. Modify yarn-site.xml with the following property:  

```
<name>yarn.scheduler.minimum-allocation-mb</name>
<value>2048</value>
```
- B. Modify yarn-sites.xml with the following property:  

```
<name>yarn.scheduler.minimum-allocation-mb</name>
<value>4096</value>
```
- C. Modify yarn-site.xml with the following property:  

```
<name>yarn.nodemanager.resource.cpu-vcores</name>
```
- D. No action is needed: YARN's dynamic resource allocation automatically optimizes the node memory and cores

**Correct Answer:** A

**QUESTION 7**

You want to node to only swap Hadoop daemon data from RAM to disk when absolutely necessary. What should you do?

- A. Delete the /dev/vmswap file on the node
- B. Delete the /etc/swap file on the node
- C. Set the ram.swap parameter to 0 in core-site.xml
- D. Set vm.swapfile file on the node
- E. Delete the /swapfile file on the node

**Correct Answer:** D

**QUESTION 8**

You are running a Hadoop cluster with a NameNode on host mynamenode. What are two ways to determine available HDFS space in your cluster?

- A. Run `hdfs fs du /` and locate the DFS Remaining value

- B. Run `hdfs dfsadmin report` and locate the DFS Remaining value
- C. Run `hdfs dfs /` and subtract NDFS Used from configured Capacity
- D. Connect to `http://mynamenode:50070/dfshealth.jsp` and locate the DFS remaining value

**Correct Answer:** B

#### QUESTION 9

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 mapreduce-reduces.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 handed 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 10

Your Hadoop cluster contains nodes in three racks. You have not configured the `dfs.hosts` property in the NameNode's configuration file. What results?

- A. The NameNode will update the `dfs.hosts` property to include machines running the DataNode daemon on the next NameNode reboot or with the command `dfsadmin refreshNodes`
- B. No new nodes can be added to the cluster until you specify them in the `dfs.hosts` file
- C. Any machine running the DataNode daemon can immediately join the cluster
- D. Presented with a blank `dfs.hosts` property, the NameNode will permit DataNodes specified in `mapred.hosts` to join the cluster

**Correct Answer:** C

#### QUESTION 11

You are running a Hadoop cluster with MapReduce version 2 (MRv2) on YARN. You consistently see that MapReduce map tasks on your cluster are running slowly because of excessive garbage collection of JVM, how do you increase JVM heap size property to 3GB to optimize performance?

- A. `yarn.application.child.java.opts=-Xsx3072m`
- B. `yarn.application.child.java.opts=-Xmx3072m`
- C. `mapreduce.map.java.opts=-Xms3072m`
- D. `mapreduce.map.java.opts=-Xmx3072m`

**Correct Answer:** C

#### QUESTION 12

You have a cluster running with a FIFO scheduler enabled. You submit a large job A to the cluster, which you expect to run for one hour. Then, you submit job B to the cluster, which you expect to run a couple of

minutes only.

You submit both jobs with the same priority.

Which two best describes how FIFO Scheduler arbitrates the cluster resources for job and its tasks?  
(Choose two)

- A. Because there is a more than a single job on the cluster, the FIFO Scheduler will enforce a limit on the percentage of resources allocated to a particular job at any given time
- B. Tasks are scheduled on the order of their job submission
- C. The order of execution of job may vary
- D. Given job A and submitted in that order, all tasks from job A are guaranteed to finish before all tasks from job B
- E. The FIFO Scheduler will give, on average, and equal share of the cluster resources over the job lifecycle
- F. The FIFO Scheduler will pass an exception back to the client when Job B is submitted, since all slots on the cluster are use

**Correct Answer:** AD

### QUESTION 13

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 14

Your company stores user profile records in an OLTP databases. You want to join these records with web server logs you have already ingested into the Hadoop file system. What is the best way to obtain and ingest these user records?

- A. Ingest with Hadoop streaming
- B. Ingest using Hive's IQAD DATA command
- C. Ingest with sqoop import
- D. Ingest with Pig's LOAD command
- E. Ingest using the HDFS put command

**Correct Answer:** C

### QUESTION 15

Which two are features of Hadoop's rack topology? (Choose two)

- A. Configuration of rack awareness is accomplished using a configuration file. You cannot use a rack

topology script.

- B. Hadoop gives preference to intra-rack data transfer in order to conserve bandwidth
- C. Rack location is considered in the HDFS block placement policy
- D. HDFS is rack aware but MapReduce daemon are not
- E. Even for small clusters on a single rack, configuring rack awareness will improve performance

**Correct Answer:** BC

To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

# Trying our product !


- ★ **100%** Guaranteed Success
- ★ **100%** Money Back Guarantee
- ★ **365 Days** Free Update
- ★ **Instant Download** After Purchase
- ★ **24x7** Customer Support
- ★ Average **99.9%** Success Rate
- ★ More than **69,000** Satisfied Customers Worldwide
- ★ Multi-Platform capabilities - **Windows, Mac, Android, iPhone, iPod, iPad, Kindle**

## Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 <b>One Year Free Update</b> <p>Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <b>Money Back Guarantee</b> <p>To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <b>Security &amp; Privacy</b> <p>We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information &amp; peace of mind.</p>
---	---	--

[Guarantee & Policy](#) | [Privacy & Policy](#) | [Terms & Conditions](#)

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © 2004-2017, All Rights Reserved.