



CCA-505^{Q&As}

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

You have installed a cluster running HDFS and MapReduce version 2 (MRv2) on YARN. You have no `dfs.hosts` entry(ies) in your `hdfs-alte.xml` configuration file. You configure a new worker node by setting `fs.default.name` in its configuration files to point to the NameNode on your cluster, and you start the DataNode daemon on that worker node.

What do you have to do on the cluster to allow the worker node to join, and start storing HDFS blocks?

- A. Nothing; the worker node will automatically join the cluster when the DataNode daemon is started.
- B. Without creating a `dfs.hosts` file or making any entries, run the command `hadoop dfsadmin refreshHadoop` on the NameNode
- C. Create a `dfs.hosts` file on the NameNode, add the worker node's name to it, then issue the command `hadoop dfsadmin refreshNodes` on the NameNode
- D. Restart the NameNode

Correct Answer: B

QUESTION 2

A slave node in your cluster has four 2TB hard drives installed (4 x 2TB). The DataNode is configured to store HDFS blocks on the disks. You set the value of the `dfs.datanode.du.reserved` parameter to 100GB. How does this alter HDFS block storage?

- A. A maximum of 100 GB on each hard drive may be used to store HDFS blocks
- B. All hard drives may be used to store HDFS blocks as long as at least 100 GB in total is available on the node
- C. 100 GB on each hard drive may not be used to store HDFS blocks
- D. 25 GB on each hard drive may not be used to store HDFS blocks

Correct Answer: B

QUESTION 3

You are configuring your cluster to run HDFS and MapReduce v2 (MRv2) on YARN. Which daemons need to be installed on your clusters master nodes? (Choose Two)

- A. ResourceManager
- B. DataNode
- C. NameNode
- D. JobTracker
- E. TaskTracker



F. HMaster

Correct Answer: AC

QUESTION 4

Given:

```
[user1@host1 ~] $ yarn application -list  
Total Applications: 3
```

Application-Id	Application-Name	Application-Type	User	Queue	State	Final State	Process	Tracking-URL
Application_1374638600275_0109	Sleep job	MAPREDUCE	user1	default	Killed	KILLED	100%	host1:54059
Application_1374638600275_0121	sleep job	MAPREDUCE	user1	default	SUCCEEDED	SUCCEEDED	100%	host1:19888/jobhistory/job/job_1374638600275_0121
Application_1374638600275_0020	sleep job	MAPREDUCE	user1	default	SUCCEEDED	SUCCEEDED	100%	host1:19888/jobhistory/job/job_1374638600275_0020

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

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

Correct Answer: A

QUESTION 5

Identify two features/issues that YARN is designed to address:

- A. Standardize on a single MapReduce API
- B. Single point of failure in the NameNode
- C. Reduce complexity of the MapReduce APIs
- D. Resource pressures on the JobTracker
- E. Ability to run frameworks other than MapReduce, such as MPI
- F. HDFS latency

Correct Answer: DE

[Latest CCA-505 Dumps](#)

[CCA-505 Practice Test](#)

[CCA-505 Exam Questions](#)