

CCA-505^{Q&As}

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

Pass Cloudera CCA-505 Exam with 100% Guarantee

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

<https://www.pass2lead.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 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?

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

Correct Answer: C

QUESTION 2

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 3

You observe that the number of spilled records from Map tasks far exceeds the number of map output records. Your child heap size is 1GB and your io.sort.mb value is set to 100 MB. How would you tune your io.sort.mb value to achieve maximum memory to disk I/O ratio?

- A. Decrease the io.sort.mb value to 0
- B. Increase the io.sort.mb to 1GB
- C. For 1GB child heap size an io.sort.mb of 128 MB will always maximize memory to disk I/O

D. Tune the io.sort.mb value until you observe that the number of spilled records equals (or is as close to equals) the number of map output records

Correct Answer: D

QUESTION 4

In CDH4 and later, which file contains a serialized form of all the directory and files inodes in the filesystem, giving the NameNode a persistent checkpoint of the filesystem metadata?

- A. fstime
- B. VERSION
- C. Fsimage_N (Where N reflects all transactions up to transaction ID N)
- D. Edits_N-M (Where N-M specifies transactions between transactions ID N and transaction ID N)

Correct Answer: C

QUESTION 5

Assuming a cluster running HDFS, MapReduce version 2 (MRv2) on YARN with all settings at their default, what do you need to do when adding a new slave node to a cluster?

- A. Nothing, other than ensuring that DNS (or /etc/hosts files on all machines) contains an entry for the new node.
- B. Restart the NameNode and ResourceManager daemons and resubmit any running jobs
- C. Increase the value of dfs.number.of.replicas in hdfs-site.xml
- D. Add a new entry to /etc/nodes on the NameNode host.
- E. Restart the NameNode daemon.

Correct Answer: B

[Latest CCA-505 Dumps](#)

[CCA-505 PDF Dumps](#)

[CCA-505 Study Guide](#)