



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.pass4itsure.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 decide to create a cluster which runs HDFS in High Availability mode with automatic failover, using Quorum-based Storage. What is the purpose of ZooKeeper in such a configuration?

- A. It manages the Edits file, which is a log changes to the HDFS filesystem.
- B. It monitors an NFS mount point and reports if the mount point disappears
- C. It both keeps track of which NameNode is Active at any given time, and manages the Edits file, which is a log of changes to the HDFS filesystem
- D. It only keeps track of which NameNode is Active at any given time
- E. Clients connect to ZoneKeeper to determine which NameNode is Active

Correct Answer: D

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 have a Hadoop cluster running HDFS, and a gateway machine external to the cluster from which clients submit jobs. What do you need to do in order to run on the cluster and submit jobs from the command line of the gateway machine?

- A. Install the implad daemon, statestored daemon, and catalogd daemon on each machine in the cluster and on the gateway node
- B. Install the impalad daemon on each machine in the cluster, the statestored daemon and catalogd daemon on one



machine in the cluster, and the impala shell on your gateway machine

C. Install the impalad daemon and the impala shell on your gateway machine, and the statestored daemon and catalog daemon on one of the nodes in the cluster

D. Install the impalad daemon, the statestored daemon, the catalogd daemon, and the impala shell on your gateway machine

E. Install the impalad daemon, statestored daemon, and catalogd daemon on each machine in the cluster, and the impala shell on your gateway machine

Correct Answer: B

QUESTION 4

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 5

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