

# 我要—PASS



質 量 更 高      務 服 更 好

一年免費升級服務

<http://www.51-pass.com>

**Exam : CCA-505**

**Title :** Cloudera Certified  
Administrator for Apache  
Hadoop (CCA-H) CDH5  
Upgrade Exam

**Version : DEMO**

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

**Answer: D**

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 rmdadmin -refreshQueue`
- C. `Yarn application -refreshJobHistory`
- D. `Yarn rmdadmin -kill application_1374638600275_0109`

**Answer: A**

Reference:[http://docs.hortonworks.com/HDPDocuments/HDP2/HDP-2.1-latest/bk\\_using-apachehadoop/content/common\\_mrv2\\_commands.html](http://docs.hortonworks.com/HDPDocuments/HDP2/HDP-2.1-latest/bk_using-apachehadoop/content/common_mrv2_commands.html)

3.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.

**Answer: B**

4.You have a 20 node Hadoop cluster, with 18 slave nodes and 2 master nodes running HDFS High Availability (HA). You want to minimize the chance of data loss in your cluster. What should you do?

- A. Add another master node to increase the number of nodes running the JournalNode which increases the number of machines available to HA to create a quorum
- B. Configure the cluster's disk drives with an appropriate fault tolerant RAID level
- C. Run the ResourceManager on a different master from the NameNode in order to load share HDFS metadata processing
- D. Run a Secondary NameNode on a different master from the NameNode in order to load provide automatic recovery from a NameNode failure

E. Set an HDFS replication factor that provides data redundancy, protecting against failure

**Answer: C**

5.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 toZoneKeeper to determine which NameNode is Active

**Answer: D**

Reference:<http://www.cloudera.com/content/cloudera-content/clouderadocs/CDH4/latest/PDF/CDH4-High-Availability-Guide.pdf>(page 15)