

CCA-505 - Cloudera Certified Administrator for Apache Hadoop (CCAH) CDH5 Upgrade Exam

<http://www.certleader.com/CCA-505-dumps.html>



1. You want to understand more about how users browse your public website. For example, you want to know which pages they visit prior to placing an order. You have a server farm of 200 web servers hosting your website. Which is the most efficient process to gather these web server logs into your Hadoop cluster for analysis?

- A. Sample the web server logs from web servers and copy them into HDFS using curl
- B. Ingest the server web logs into HDFS using Flume
- C. Import all users clicks from your OLTP databases into Hadoop using Sqoop
- D. Write a MapReduce job with the web servers from mappers and the Hadoop cluster nodes reducers
- E. Channel these clickstream into Hadoop using Hadoop Streaming

Answer: A,B

2. Assume you have a file named foo.txt in your local directory. You issue the following three commands:

```
Hadoop fs -mkdir input
```

```
Hadoop fs -put foo.txt input/foo.txt
```

```
Hadoop fs -put foo.txt input
```

What happens when you issue that third command?

- A. The write succeeds, overwriting foo.txt in HDFS with no warning
- B. The write silently fails
- C. The file is uploaded and stored as a plain named input
- D. You get an error message telling you that input is not a directory
- E. You get an error message telling you that foo.txt already exists. The file is not written to HDFS
- F. You get an error message telling you that foo.txt already exists, and asking you if you would like to overwrite
- G. You get a warning that foo.txt is being overwritten

Answer: E

3. 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`

Answer: C

4. 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 `impslad` 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

Answer: B

5. Your cluster is configured with HDFS and MapReduce version 2 (MRv2) on YARN. What is the result when you execute: `hadoop jar samplejar.jar MyClass` on a client machine?

- A. `SampleJar.jar` is sent to the ApplicationMaster which allocates a container for `Sample.jar`
- B. `SampleJar.jar` is serialized into an XML file which is submitted to the ApplicationMaster
- C. `SampleJar.jar` is sent directly to the ResourceManager
- D. `SampleJar.jar` is placed in a temporary directory in HDFS

Answer: A

6. Your Hadoop cluster is configured with HDFS and MapReduce version 2 (MRv2) on YARN. Can you configure a worker node to run a NodeManager daemon but not a DataNode daemon and still have a functional cluster?

- A. Yes. The daemon will receive data from the NameNode to run Map tasks
- B. Yes. The daemon will get data from another (non-local) DataNode to run Map tasks
- C. Yes. The daemon will receive Reduce tasks only

Answer: A

7. Your cluster has the following characteristics:

- ? A rack aware topology is configured and on
- ? Replication is not set to 3
- ? Cluster block size is set to 64 MB

Which describes the file read process when a client application connects into the cluster and requests a 50MB file?

- A. The client queries the NameNode which retrieves the block from the nearest DataNode to the client and then passes that block back to the client.
- B. The client queries the NameNode for the locations of the block, and reads from a random location in the list it retrieves to eliminate network I/O leads by balancing which nodes it retrieves data from at any given time.
- C. The client queries the NameNode for the locations of the block, and reads all three copies. The first copy to complete transfer to the client is the one the client reads as part of Hadoop's speculative execution framework.
- D. The client queries the NameNode for the locations of the block, and reads from the first location in the list it receives.

Answer: A

8. 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)

Answer: C

Explanation: Reference: <http://mikepluta.com/tag/namenode/>

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your CCA-505 Exam with Our Prep Materials Via below:

<http://www.certleader.com/CCA-505-dumps.html>