

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.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 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

Correct Answer: AB

QUESTION 2

Each node in your Hadoop cluster, running YARN, has 64 GB memory and 24 cores. Your yarn-site.xml has the following configuration:

```
yarn.nodemanager.resource.memory-mb
```

```
32768
```

```
yarn.nodemanager.resource.cpu-vcores
```

```
23
```

You want YARN to launch no more than 16 containers per node. What should you do?

- A. No action is needed: YARN's dynamic resource allocation automatically optimizes the node memory and cores
- B. Modify yarn-site.xml with the following property: yarn.nodemanager.resource.cpu-vcores 16
- C. Modify yarn-site.xml with the following property: yarn.scheduler.minimum-allocation-mb 2048
- D. Modify yarn-site.xml with the following property: yarn.scheduler.minimum-allocation-mb 4096

Correct Answer: B

QUESTION 3

Your cluster implements HDFS High Availability (HA). Your two NameNodes are named nn01 and nn02. What occurs when you execute the command: `hdfs haadmin failover nn01 nn02`

- A. nn02 becomes the standby NameNode and nn01 becomes the active NameNode
- B. nn02 is fenced, and nn01 becomes the active NameNode
- C. nn01 becomes the standby NameNode and nn02 becomes the active NameNode
- D. nn01 is fenced, and nn02 becomes the active NameNode

Correct Answer: D

QUESTION 4

For each YARN Job, the Hadoop framework generates task log files. Where are Hadoop's files stored?

- A. In HDFS, In the directory of the user who generates the job
- B. On the local disk of the slave node running the task
- C. Cached In the YARN container running the task, then copied into HDFS on job completion
- D. Cached by the NodeManager managing the job containers, then written to a log directory on the NameNode

Correct Answer: B

QUESTION 5

You are running a Hadoop cluster with MapReduce version 2 (MRv2) on YARN. You consistently see that MapReduce map tasks on your cluster are running slowly because of excessive garbage collection of JVM, how do you increase JVM heap property to 3GB to optimize performance?

- A. `Yarn.application.child.java.opts=-Xax3072m`
- B. `Yarn.application.child.java.opts=-3072m`
- C. `Mapreduce.map.java.opts=-Xmx3072m`
- D. `Mapreduce.map.java.opts=-Xms3072m`

Correct Answer: C

QUESTION 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

Correct Answer: A

QUESTION 7

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 8

What processes must you do if you are running a Hadoop cluster with a single NameNode and six DataNodes, and you want to change a configuration parameter so that it affects all six DataNodes.

- A. You must modify the configuration file on each of the six DataNode machines.
- B. You must restart the NameNode daemon to apply the changes to the cluster
- C. You must restart all six DataNode daemon to apply the changes to the cluster
- D. You don't need to restart any daemon, as they will pick up changes automatically
- E. You must modify the configuration files on the NameNode only. DataNodes read their configuration from the master nodes.

Correct Answer: BE

QUESTION 9

Which YARN process runs as "controller O" of a submitted job and is responsible for resource requests?

- A. ResourceManager
- B. NodeManager

- C. JobHistoryServer
- D. ApplicationMaster
- E. JobTracker
- F. ApplicationManager

Correct Answer: D

QUESTION 10

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

[Latest CCA-505 Dumps](#)

[CCA-505 Practice Test](#)

[CCA-505 Exam Questions](#)