



CCA-500^{Q&As}

Cloudera Certified Administrator for Apache Hadoop (CCA-H)

Pass Cloudera CCA-500 Exam with 100% Guarantee

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

<https://www.geekcert.com/cca-500.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

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

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

```
32768
```

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

```
12
```

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

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

Correct Answer: A

QUESTION 2

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 analysis?

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

Correct Answer: B



QUESTION 3

Which YARN process run as "container 0" of a submitted job and is responsible for resource requests?

- A. ApplicationManager
- B. JobTracker
- C. ApplicationMaster
- D. JobHistoryServer
- E. ResourceManager
- F. NodeManager

Correct Answer: C

QUESTION 4

You have installed a cluster HDFS and MapReduce version 2 (MRv2) on YARN. You have no dfs.hosts entry(ies) in your hdfs-site.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. Without creating a dfs.hosts file or making any entries, run the commands `hadoop dfsadmin refreshNodes` on the NameNode
- B. Restart the NameNode
- C. Creating 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. Nothing; the worker node will automatically join the cluster when NameNode daemon is started

Correct Answer: A

QUESTION 5

You want to node to only swap Hadoop daemon data from RAM to disk when absolutely necessary. What should you do?

- A. Delete the `/dev/vmswap` file on the node
- B. Delete the `/etc/swap` file on the node
- C. Set the `ram.swap` parameter to 0 in `core-site.xml`
- D. Set `vm.swapfile` file on the node
- E. Delete the `/swapfile` file on the node

Correct Answer: D



VCE & PDF

GeekCert.com

<https://www.geekcert.com/cca-500.html>

2024 Latest geekcert CCA-500 PDF and VCE dumps Download

[CCA-500 PDF Dumps](#)

[CCA-500 Practice Test](#)

[CCA-500 Braindumps](#)