

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.pass2lead.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

Your company stores user profile records in an OLTP databases. You want to join these records with web server logs you have already ingested into the Hadoop file system. What is the best way to obtain and ingest these user records?

- A. Ingest with Hadoop streaming
- B. Ingest using Hive's IQAD DATA command
- C. Ingest with sqoop import
- D. Ingest with Pig's LOAD command
- E. Ingest using the HDFS put command

Correct Answer: C

QUESTION 2

Cluster Summary:

45 files and directories, 12 blocks = 57 total. Heap size is 15.31 MB/193.38MB(7%)

Configured capacity	:	17.33GB
DFS Used	:	144KB
Non DFS Used	:	5.49GB
DFS Remaining	:	11.84GB
DFS Used %	:	0%
DFS Remaining %	:	68.32GB
Live Nodes	:	6
Dead Nodes	:	1
Decommissioning Nodes	:	0
Number of Under-Replicated Blocks	:	6

Refer to the above screenshot.

You configure a Hadoop cluster with seven DataNodes and on of your monitoring UIs displays the details shown in the exhibit.

What does the this tell you?

- A. The DataNode JVM on one host is not active
- B. Because your under-replicated blocks count matches the Live Nodes, one node is dead, and your DFS Used % equals 0%, you can't be certain that your cluster has all the data you've written it.
- C. Your cluster has lost all HDFS data which had bocks stored on the dead DatNode

D. The HDFS cluster is in safe mode

Correct Answer: A

QUESTION 3

On a cluster running CDH 5.0 or above, you use the `hadoop fs put` command to write a 300MB file into a previously empty directory using an HDFS block size of 64 MB. Just after this command has finished writing 200 MB of this file, what would another user see when they look in the directory?

A. The directory will appear to be empty until the entire file write is completed on the cluster

B. They will see the file with a `._COPYING_` extension on its name. If they view the file, they will see contents of the file up to the last completed block (as each 64MB block is written, that block becomes available)

C. They will see the file with a `._COPYING_` extension on its name. If they attempt to view the file, they will get a `ConcurrentFileAccessException` until the entire file write is completed on the cluster

D. They will see the file with its original name. If they attempt to view the file, they will get a `ConcurrentFileAccessException` until the entire file write is completed on the cluster

Correct Answer: B

QUESTION 4

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 5

What two 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. (Choose two)

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

Correct Answer: AD

[Latest CCA-500 Dumps](#)

[CCA-500 Study Guide](#)

[CCA-500 Exam Questions](#)