

CCA-505^{Q&As}

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

Pass Cloudera CCA-505 Exam with 100% Guarantee

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

<https://www.leads4pass.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 a node to only swap Hadoop daemon data from RAM to disk when absolutely necessary. What should you do?

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

Correct Answer: B

QUESTION 2

A slave node in your cluster has four 2TB hard drives installed (4 x 2TB). The DataNode is configured to store HDFS blocks on the disks. You set the value of the dfs.datanode.du.reserved parameter to 100GB. How does this alter HDFS block storage?

- A. A maximum of 100 GB on each hard drive may be used to store HDFS blocks
- B. All hard drives may be used to store HDFS blocks as long as at least 100 GB in total is available on the node
- C. 100 GB on each hard drive may not be used to store HDFS blocks
- D. 25 GB on each hard drive may not be used to store HDFS blocks

Correct Answer: B

QUESTION 3

Which process instantiates user code, and executes map and reduce tasks on a cluster running MapReduce V2 (MRv2) on YARN?

- A. NodeManager
- B. ApplicationMaster
- C. ResourceManager
- D. TaskTracker
- E. JobTracker
- F. DataNode
- G. NameNode

Correct Answer: E

QUESTION 4

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)

Correct Answer: C

QUESTION 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 to ZooKeeper to determine which NameNode is Active

Correct Answer: D

QUESTION 6

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 of 64MB. Just after this command has finished writing 200MB of this file, what would another user see when they look in the directory?

- A. 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
- B. 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.
- C. 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)
- D. The directory will appear to be empty until the entire file write is completed on the cluster

Correct Answer: C

QUESTION 7

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.

Correct Answer: B

QUESTION 8

Which Yarn daemon or service monitors a Container's per-application resource usage (e.g, memory, CPU)?

- A. NodeManager
- B. ApplicationMaster
- C. ApplicationManagerService
- D. ResourceManager

Correct Answer: A

QUESTION 9

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 10

On a cluster running MapReduce v2 (MRv2) on YARN, a MapReduce job is given a directory of 10 plain text as its input directory. Each file is made up of 3 HDFS blocks. How many Mappers will run?

- A. We cannot say; the number of Mappers is determined by the ResourceManager
- B. We cannot say; the number of Mappers is determined by the ApplicationManager
- C. We cannot say; the number of Mappers is determined by the developer
- D. 30
- E. 3
- F. 10

Correct Answer: E

[Latest CCA-505 Dumps](#)

[CCA-505 VCE Dumps](#)

[CCA-505 Braindumps](#)